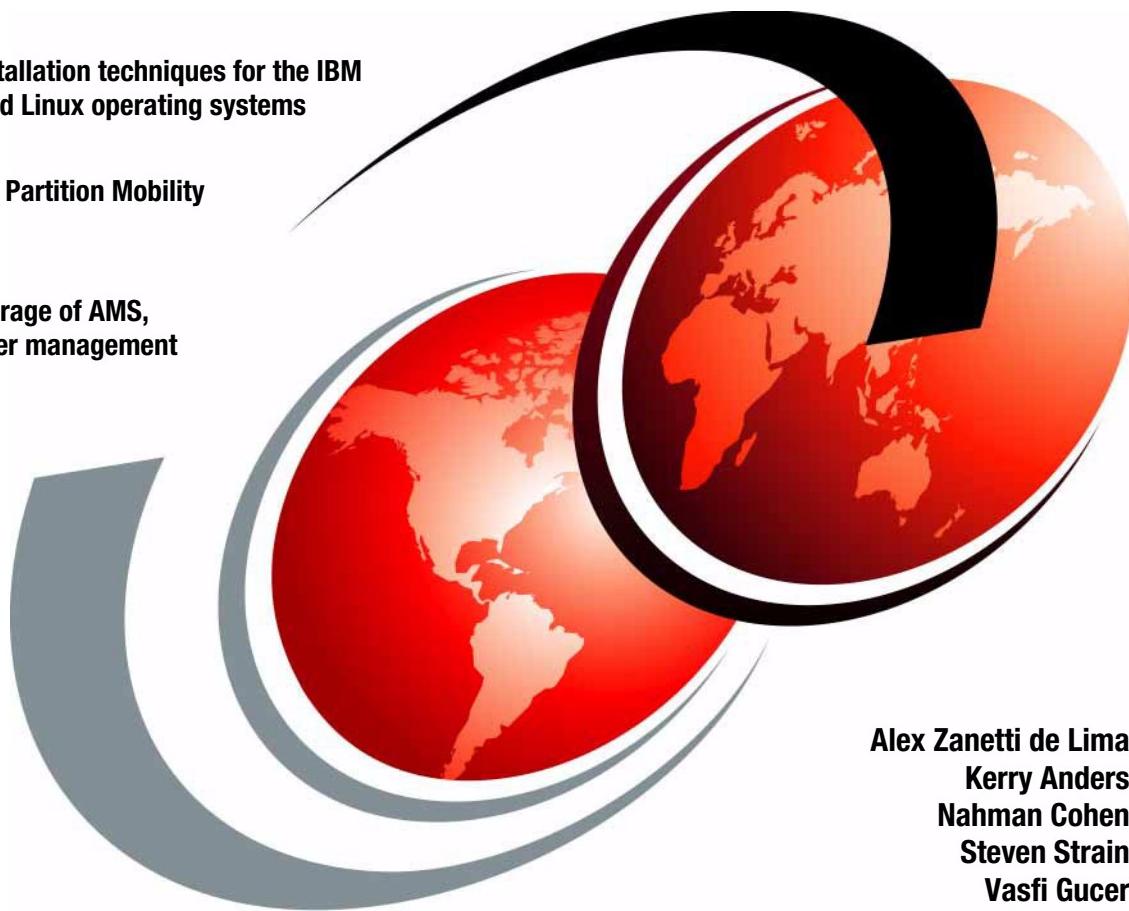


IBM BladeCenter JS23 and JS43 Implementation Guide

Featuring installation techniques for the IBM AIX, IBM i, and Linux operating systems

Showing Live Partition Mobility scenarios

Detailed coverage of AMS, IVM and power management



Alex Zanetti de Lima
Kerry Anders
Nahman Cohen
Steven Strain
Vasfi Gucer

Redbooks



International Technical Support Organization

**IBM BladeCenter JS23 and JS43 Implementation
Guide**

May 2009

Note: Before using this information and the product it supports, read the information in "Notices" on page xxv.

First Edition (May 2009)

This edition applies to IBM BladeCenter JS23, IBM BladeCenter JS43, IBM AIX Version 6.1, IBM i 6.1, Red Hat Enterprise Linux for POWER Version 5.3, SUSE Linux Enterprise Server 11 for POWER..

© Copyright International Business Machines Corporation 2009. All rights reserved.

Note to U.S. Government Users Restricted Rights -- Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

Figuresxi
Tablesxxiii
Noticesxxv
Trademarksxxvi
Prefacexxix
The team that wrote this bookxxix
Acknowledgementsxxx
Become a published authorxxxi
Comments welcome.....	.xxxi
Part 1. The IBM BladeCenter JS23 and JS43 servers1
Chapter 1. Introduction to IBM BladeCenter3
1.1 Highlights of BladeCenter4
1.2 IBM BladeCenter is the right choice, open, easy and green6
Chapter 2. General description9
2.1 Overview of the JS23 and JS43 Express blade servers10
2.2 Features and specifications.....	.10
2.3 Blade Server features16
2.4 Physical specifications BladeCenter JS2318
2.4.1 Minimal and optional features of BladeCenter JS23.....	.18
2.4.2 Processor features20
2.4.3 Memory features20
2.4.4 Memory DIMMs.....	.20
2.4.5 Internal disk.....	.21
2.5 Physical specifications BladeCenter JS4321
2.5.1 Minimal and optional features of BladeCenter JS43.....	.21
2.5.2 Processor features22
2.5.3 Memory features23
2.5.4 Memory DIMMs.....	.23
2.5.5 Internal disk.....	.23
2.6 IBM BladeCenter chassis24
2.6.1 BladeCenter H.....	.25
2.6.2 BladeCenter S.....	.29
2.6.3 BladeCenter HT33

2.6.4 Number of IBM BladeCenter JS23 and JS43 Express in Supported Blade Center Chassis	37
2.6.5 IBM Director	38
Chapter 3. Technical description of the hardware architecture	41
3.1 POWER6 processor	43
3.1.1 Decimal floating point	44
3.2 Simultaneous Multi Threading (SMT)	45
3.3 AltiVec (SIMD)	46
3.4 IBM EnergyScale technology	47
3.4.1 Thermal Power Management Device (TPMD)	49
3.5 Flexible Support Processor	49
3.6 Management Card	50
3.6.1 Anchor chip on Management Card	50
3.7 Memory subsystem	50
3.7.1 Memory description of IBM BladeCenter JS23 and JS43 Express	51
3.7.2 Memory placement rules for IBM BladeCenter JS23 and JS43	51
3.7.3 Memory RAS	53
3.8 I/O subsystem	53
3.8.1 PCI express	53
3.8.2 I/O Expansion Cards	54
3.8.3 Universal Serial Bus (USB) subsystem	55
3.8.4 Integrated Virtual Ethernet (IVE)	56
3.8.5 Integrated Video Card	56
3.8.6 Serial Attached SCSI (SAS) storage subsystem	57
3.9 PowerVM	61
3.10 Operating system support	62
3.10.1 AIX	62
3.10.2 Linux	63
3.10.3 IBM i	64
3.11 Systems management	64
3.11.1 BladeCenter Advanced Management Web interface	64
3.11.2 IBM Director	65
3.11.3 Cluster Systems Management (CSM)	66
Part 2. System configuration and management	69
Chapter 4. System planning and configuration using VIOS with IVM	71
4.1 Planning considerations	73
4.1.1 General considerations	73
4.1.2 Internal and external storage considerations	76
4.2 VIOS system management using IVM	83
4.2.1 VIOS installation considerations	84
4.2.2 IVM user interface	84

4.2.3 VIOS/IVM command line interface	85
4.3 First VIOS login	86
4.3.1 Password set	86
4.3.2 License acceptance	86
4.3.3 Initial network setup	87
4.4 First IVM connection	93
4.4.1 Connecting to IVM	93
4.4.2 Verify and modify VIOS partition memory and processors	95
4.5 VIOS network management and setup	100
4.5.1 Host Ethernet Adapters	100
4.5.2 Virtual Ethernet Adapters and SEA	103
4.5.3 Physical adapters	109
4.5.4 VLANs and virtual networks	113
4.6 VIOS Storage Configuration and Management	121
4.6.1 Physical volumes	122
4.6.2 Storage pools	124
4.6.3 Virtual disks	127
4.6.4 Optical and Tape devices	132
4.7 Partition configuration for Virtual I/O Client (VIOC)	144
4.7.1 Live Partition Mobility considerations	144
4.7.2 Partition name and environment	145
4.7.3 Partition name	146
4.7.4 Partition Memory	147
4.7.5 Partition Processors	148
4.7.6 Partition Ethernet	149
4.7.7 Partition Storage Types and Assignments	150
4.7.8 Optical and tape devices	152
4.7.9 Physical adapters	153
4.7.10 Partition Summary	154
4.7.11 Partition properties changes and DLPAR operations	156
4.8 Console access and activating a partition	166
4.8.1 Opening a virtual terminal	167
4.8.2 Activating a partition	168
4.8.3 Shutdown a VIO Client partition	172
4.8.4 Shutdown the VIO Server	175
Chapter 5. Active Memory Sharing configuration using IVM	177
5.1 Planning considerations	179
5.1.1 Prerequisites	179
5.1.2 Workload	180
5.1.3 Paging devices	181
5.1.4 Collaborative Memory Manager	181
5.1.5 CPU resource sizing	181

5.1.6 Memory weight	182
5.1.7 Consolidation factors.	182
5.2 Configuration of Active Memory Sharing using IVM	183
5.2.1 Defining the shared memory pool and paging storage pool	183
5.2.2 Creating dedicated paging devices for partitions	191
5.2.3 Creating shared memory LPARs	194
5.2.4 Shared memory partitions and dedicated paging devices	204
5.2.5 Active Memory Sharing DLPAR operations	208
5.3 Active Memory Sharing summary	209
Chapter 6. IBM AIX V6.1 installation	211
6.1 Install IBM AIX 6.1 in a PowerVM client partition	212
6.1.1 Create a virtual media library	212
6.1.2 Prepare the PowerVM client partition	218
6.1.3 Install AIX 6.1 in a logical partition of the Virtual IO Server.	231
Chapter 7. IBM i V6.1 installation	241
7.1 Preparing for installation	242
7.1.1 Software installation process	242
7.1.2 Hardware environments	244
7.1.3 BladeCenter hardware preparation.	247
7.1.4 VIO Server software environments	251
7.1.5 Network considerations.	251
7.1.6 Storage consideration BladeCenter H.	252
7.1.7 Disk consideration in BladeCenter S.	252
7.1.8 Disk configuration in BladeCenter S	254
7.1.9 Individual BladeCenter S disk configuration	259
7.2 IBM System Access for Windows V6R1	260
7.2.1 Preparation System i Access for Windows	260
7.2.2 Installation System i Access for Windows.	261
7.3 Creating an IBM i 6.1 partition.	271
7.3.1 IBM i 6.1 minimum requirements	271
7.3.2 VIO Server configuration.	272
7.3.3 Creating an IBM i partition.	272
7.3.4 Partition properties	277
7.3.5 IBM i 6.1 install media preparation	280
7.3.6 Connect the System i LAN console.	280
7.3.7 IBM i 6.1 IPL types	292
7.4 Installing the IBM i 6.1 Operating System	303
7.4.1 IBM i 6.1 installing PTFs	311
7.4.2 IBM i 6.1 TCP/IP setup	312
7.4.3 IBM i 6.1 Navigator for i.	312
7.5 IBM i 6.1 Backup/Restore	314

7.5.1	Creating a virtual media library for backup	316
7.5.2	Creating Virtual Media Library using IVM	322
7.5.3	Adding Image Files to Media Library.	324
7.5.4	Attaching a remote PC file or Media device	327
7.5.5	IBM Tivoli Storage Manager	331
7.5.6	IBM i 6.1 shutdown and restart	332
Chapter 8. Red Hat Enterprise V5.3 Linux installation	335
8.1	Supported Red Hat operating systems	336
8.1.1	Considerations and prerequisites	336
8.2	Linux LPAR installation using DVD	337
8.3	Linux network installation (detailed)	341
8.4	Native Red Hat Enterprise Linux 5.3 installation.	353
8.5	Red Hat Enterprise Linux 5.3 automated installation	354
8.5.1	Kickstart profile creation methods.	354
8.5.2	Create Kickstart file using Kickstart Configurator	356
8.5.3	Performing a Kickstart installation.	364
8.6	IBM service and productivity tools.	365
Chapter 9. SUSE Linux Enterprise Server V11 installation	367
9.1	Supported operating systems	368
9.2	Considerations and prerequisites	368
9.2.1	PowerVM LPAR considerations	368
9.2.2	SUSE Linux Enterprise Server 11 installation considerations.	368
9.3	Linux LPAR installation using DVD	369
9.4	Linux network installation (detailed)	374
9.5	Native SLES 11 installation.	394
9.6	SLES 11 automated installation	395
9.7	IBM service and productivity tools.	396
Chapter 10. JS23 and JS43 power management using EnergyScale technology	397
10.1	Power management through the AMM	398
10.1.1	Using the AMM Web UI for blade power management.	398
10.1.2	Using the AMM CLI UI for blade power management.	404
10.2	Power management through AEM	407
10.2.1	Active Energy Manager Console.	408
10.2.2	AEM Energy Properties.	410
10.2.3	BladeCenter Energy Properties	413
10.2.4	Creating Power Policies	421
Chapter 11. Performing Live Partition Mobility	433
11.1	Requirements	434
11.1.1	Hardware.	434

11.1.2 Firmware	434
11.1.3 VIOS version	434
11.1.4 PowerVM Enterprise	436
11.1.5 LPAR OS versions	437
11.2 Preparation	438
11.2.1 VIOS (source and target) requirements	438
11.2.2 Networking	442
11.2.3 Partition requirements	443
11.3 Migrating the LPAR	448
11.3.1 Using the IVM UI	448
11.3.2 From the command line	456
Chapter 12. System maintenance and diagnostics	459
12.1 Firmware updates	460
12.1.1 Commit new firmware to TEMP side	460
12.1.2 Starting the firmware image from the TEMP side	462
12.1.3 Verify current firmware level	469
12.1.4 Interpreting the firmware file name	470
12.1.5 Get the latest firmware image from IBM	471
12.2 System diagnostics	472
12.2.1 Diagnostic tools	472
12.2.2 Reference codes	479
12.2.3 Diagnostics tasks through the Advanced Management Module	480
12.2.4 IBM i partition diagnostics and errors	483
Part 3. Appendixes	491
Appendix A. Consoles, SMS, and Open Firmware	493
Consoles of the IBM BladeCenter JS23 and JS43	494
Graphical console	494
Serial Over LAN	500
System Management Services menu	501
Power on the IBM BladeCenter JS23 or JS43 using the Advanced Management Module Web interface	502
Power on the IBM BladeCenter JS23 and JS43 using a Telnet or SSH session into the Advanced Management Module	503
Choosing the active console after first power on	506
IBM BladeCenter JS23 and JS43 Power On Self Test (POST)	507
Select the active console using the System Maintenance Services menu	508
Open Firmware interface	509
Get access to the firmware prompt	510
QLogic host bus adapter	512
Emulex host bus adapter	516

Appendix B. SUSE Linux Enterprise Server AutoYaST	521
AutoYaST introduction	522
AutoYaST profile creation methods	522
Create an AutoYaST profile using YaST Control Center	522
Appendix C. Additional Linux installation configuration options	535
Basic preparations for a Linux network installation	536
Installing Linux using the network - General remarks	536
Configuring a BOOTP or DHCP service	537
Configuring a Trivial File Transfer Protocol service	539
Configuring a Network File System Protocol service	541
Virtual optical device setup and installation	544
Appendix D. Service and productivity tools for Linux	545
IBM service diagnostic aids and productivity tools overview	546
Install tools on Red Hat Enterprise Linux 5/SLES 11 running on BladeCenter servers	548
Install tools on Red Hat Enterprise Linux 5/SLES 11 running on IVM-managed servers	551
Abbreviations and acronyms	557
Related publications	563
IBM Redbooks	563
Online resources	563
How to get Redbooks	568
Help from IBM	568
Index	571

Figures

1-1 With and without blade servers	4
2-1 JS23 Blade physical layout	19
2-2 JS43 Multiple Expansion Unit (MPE)	22
2-3 Front view of BladeCenter H	26
2-4 Rear view of BladeCenter H	27
2-5 BladeCenter S front view	30
2-6 BladeCenter S rear view	31
2-7 IBM BladeCenter HT front view	34
2-8 IBM BladeCenter HT rear view	35
2-9 Maximum number of IBM BladeCenter JS23 Express blades per IBM Blade Center chassis	37
2-10 Maximum number of IBM BladeCenter JS23 and JS43 Express per IBM Blade Center chassis	38
3-1 POWER6 processor module (4-way DMC)	44
3-2 IBM BladeCenter JS23 Express memory wiring layout (Base planar - P1) .	51
3-3 Recommended memory placement for IBM BladeCenter JS23 and JS43 Express	52
3-4 JS23 Base planar with a CFFh High Speed Daughter Card and CIOv Card and 2"5 HDD	54
3-5 IBM BladeCenter JS23 and JS43 Express BCS SAS Topology	60
3-6 IBM BladeCenter JS23 and JS43 Express BCH and BCHT SAS Topology .	61
4-1 Expansion card port path through BladeCenter S mid plane	74
4-2 Expansion card port path through BladeCenter H mid plane	75
4-3 Active SAS Pass-through Expansion Card CIOv	77
4-4 Qlogic Ethernet and 4 Gb Fibre Channel "combo card" CFFhTh	78
4-5 Multi-Switch Interconnect Module	79
4-6 IVM navigation and work areas	85
4-7 The Welcome window	94
4-8 Guided Setup window	95
4-9 View/Modify Partitions window	96
4-10 Partition Properties, General tab	97
4-11 Partition Properties, Memory tab	98
4-12 Partition Properties, Processing tab	99
4-13 View/Modify Host Ethernet Adapters window	101
4-14 HEA Port Properties	102
4-15 HEA Port Properties, Connected Partitions	103
4-16 HEA port setting for Ethernet bridging	104

4-17	View/Modify Virtual Ethernet showing Initialize Virtual Ethernet option	105
4-18	View/Modify Virtual Ethernet window	106
4-19	View/Modify Virtual Ethernet Bridge tab	107
4-20	Physical adapter selection for SEA creation	108
4-21	Successful SEA creation result	109
4-22	View/Modify Physical Adapters window	110
4-23	Physical Adapter assignment to new partition	111
4-24	View/Modify Physical Adapter window showing change of ownership of Gigabit Ethernet Adapter	112
4-25	Default VIO Server virtual Ethernet Adapters shown by IVM	115
4-26	View Modify Virtual Ethernet view showing new virtual adapter with multiple VLANs	117
4-27	Creating a SEA using an IEEE 802.1q enabled virtual adapter	118
4-28	VIOC adapter to VIOS virtual Ethernet mapping	120
4-29	View and modify virtual storage	122
4-30	Physical volumes shown in IVM	123
4-31	Storage pools shown in IVM	125
4-32	Create new storage pool	126
4-33	Newly created storage pool shown in IVM	126
4-34	Reduce or delete a storage pool	127
4-35	Delete storage pool	127
4-36	Create virtual disk	128
4-37	Virtual disk settings	129
4-38	The newly created virtual disk	129
4-39	Extend virtual disk	130
4-40	Virtual disk extension settings	130
4-41	Extended virtual disk	131
4-42	Delete virtual disk	132
4-43	Confirm deletion of the virtual disk	132
4-44	Physical optical and tape devices in IVM	134
4-45	Change physical optical device assignment	135
4-46	Create media library	136
4-47	Media library size and storage pool	137
4-48	Add media to media library	138
4-49	Add new media - settings	139
4-50	Performing task	140
4-51	Newly created media with the copy operation in progress	141
4-52	Modify partition assignment	142
4-53	Modify media partition assignment	143
4-54	New assigned media to partitions	144
4-55	View/Modify Partition	146
4-56	Create Partition: Name	147
4-57	Create Partitions: Memory	148

4-58 Create Partition: Processors	149
4-59 Create Partition: Ethernet	150
4-60 Create Partition: Storage Type	151
4-61 Logical Partition: Storage	152
4-62 Create Partition: Optical	153
4-63 Create Partition: Physical Adapters	154
4-64 Create Partition: Summary	155
4-65 View/Modify Partition showing new partition.....	156
4-66 Partition Properties General tab	157
4-67 DLPAR retrieved capabilities	158
4-68 Partition Properties Memory tab	159
4-69 Partition Properties, Processing tab for shared pool.....	160
4-70 Partition Properties, Processing tab for dedicated processors	161
4-71 Partition Properties, Ethernet tab	162
4-72 Partition Properties, Storage tab.....	163
4-73 Partition Properties Optical/Tape Devices tab	164
4-74 Partition Properties, changing the current mounted media.....	165
4-75 Partition Properties, Physical Adapters tab.....	166
4-76 Start virtual terminal for an LPAR	167
4-77 Virtual Terminal started from IVM UI.....	168
4-78 Activating a partition	169
4-79 Confirm partition activation	170
4-80 Shutdown an LPAR.....	173
4-81 Partition shutdown options	174
4-82 VIOS partition shutdown options.....	175
5-1 Start Active Memory Sharing configuration with View/Modify Virtual Storage 184	
5-2 Starting the storage pool creation	185
5-3 Naming the storage pool and selecting backing devices	186
5-4 Storage pool list with new pool for paging added	187
5-5 Defining a shared memory pool	188
5-6 Shared memory pool configuration values	189
5-7 Shared memory pool defined state	190
5-8 Shared memory pool information after creation	191
5-9 Creating dedicated paging devices for LPARS using shared memory ..	192
5-10 Dedicated device selection for share memory use	193
5-11 Dedicated device defined to paging space devices	194
5-12 Creating a shared memory partition name	195
5-13 Selecting memory mode and amount for a shared memory partition ..	196
5-14 Selecting the number of processors in a shared memory partition ..	197
5-15 Ethernet selection for a shared memory partition	198
5-16 Storage selection for a shared memory partition	199
5-17 Storage selection for a shared memory partition	200

5-18 Optical and tape selections for a shared memory partition	201
5-19 Summary of selections for a shared memory partition	202
5-20 View/Modify Partition window showing newly created shared memory partition	203
5-21 shared memory pool with paging space assignments in paging pool	204
5-22 Shared memory pool view showing both types of paging devices	205
5-23 Shared memory pool view showing assigned dedicated paging device	206
5-24 Partition memory properties showing maximum memory and paging space changes	207
5-25 shared memory pool after partition maximum memory values changed	208
6-1 Media library - create storage pool	212
6-2 Media library - select the physical volume	213
6-3 Media library - new storage pool	213
6-4 Media library - create library	214
6-5 Media library - specify library size and volume group	214
6-6 Media library - add media	215
6-7 Error message using a wrong media name	215
6-8 Media library - specify details about media creation	216
6-9 Media library - performing task	217
6-10 Media library - newly created media	217
6-11 IVM login window	218
6-12 View/Modify Partitions page after logon	219
6-13 View/Modify Storage	220
6-14 Available physical volumes	221
6-15 View/ Modify Partition - Create Partition	221
6-16 Create partition - define name	222
6-17 Create partition - configure the amount of memory	223
6-18 Create partition - CPU configuration	224
6-19 Create partition - ethernet network	225
6-20 Create partition - storage type	226
6-21 Create partition - select physical volumes	227
6-22 Create partition - modify virtual optical device	228
6-23 Create partition - modify current media of virtual optical device	229
6-24 Create partition - virtual optical device	229
6-25 Create partition - summary	230
6-26 Newly created AIX/Linux partition	231
6-27 Activate a partition	232
6-28 Confirm partition activation	232
6-29 Open a virtual terminal to the partition	233
6-30 Virtual terminal connection	234
7-1 IBM i 6.1 installation process	243
7-2 Firmware information and download	248
7-3 Example: Partial list of available downloads by type	250

7-4	SAS Connection module login	254
7-5	SAS connection module welcome.	255
7-6	SAS connection module zone groups.	256
7-7	SAS connection module zoning	257
7-8	AMM SAS configuration zone	258
7-9	AMM SAS configuration zone 9	259
7-10	Setup LAN Console.	261
7-11	IBM System i Access for Windows welcome screen.	262
7-12	IBM System i Access for Windows Licence Agreement	263
7-13	IBM System i Access for Windows install location	264
7-14	IBM System i Access for Windows Primary language	265
7-15	IBM System i Access for Windows Setup Type	266
7-16	IBM System i Access for Windows Restricted Features	267
7-17	IBM System i Access for Windows installation progress.	268
7-18	IBM System i Access for Windows installation completed	269
7-19	IBM System i Access for Windows Reboot.	269
7-20	IBM System i Access for Windows Welcome	270
7-21	View Modify Partitions - Create Partition option	272
7-22	Partition id, name and environment options	273
7-23	Partition memory definition panel	274
7-24	Partition processor selection	274
7-25	Partition ethernet selection	275
7-26	Select storage type	275
7-27	Disk selection	276
7-28	Partition optical selections.	277
7-29	Load Source and Console Identification fields	278
7-30	Partition memory allocation.	279
7-31	Partition processing properties tab	279
7-32	Processing units value change	280
7-33	IBM System i Operations Console	281
7-34	IBM System I operator console	281
7-35	IBM System i Operations Console Welcome	282
7-36	IBM System i Operations Console - choose a configuration.	283
7-37	IBM System i Operations Console - enter the Service host name	284
7-38	IBM System i operator console	285
7-39	IBM System i Operations Console - enter the Service TCP/IP Address	286
7-40	Specify Interface Information	287
7-41	System Properties - Serial number.	288
7-42	Partition ID	288
7-43	IBM System i Operations Console - enter System serial number	289
7-44	IBM System i Operations Console - enter Service tools device ID	290
7-45	IBM System i Operations Console - finalizing the setup	291
7-46	IBM System i Operations Console	291

7-47	Connect console session	292
7-48	LIC initial installation screen	293
7-49	Confirm Language setup	294
7-50	Install LIC	295
7-51	Select source disk	296
7-52	Confirm source device	297
7-53	Select options	298
7-54	Confirm definition	299
7-55	Initialize disk	300
7-56	Install LIC status	301
7-57	Attention Report	302
7-58	Install the operating system	303
7-59	Select install device	304
7-60	Confirm installation	305
7-61	Select language feature	306
7-62	LIC install confirm language	307
7-63	IPL in progress	308
7-64	Set date and time	309
7-65	Installation status	310
7-66	Sign On screen	311
7-67	IBM Systems Navigator for i	313
7-68	Physical Tape Devices	315
7-69	IVM Create Storage Pool	317
7-70	IVM Create blank media	318
7-71	IVM Virtual optical device created	319
7-72	Virtual optical device assign to partition	320
7-73	IVM Virtual optical device assignment done	321
7-74	Virtual optical device check device	322
7-75	Create Storage Pool option	323
7-76	Selecting storage pool name, size and resource	323
7-77	Create Media Library	324
7-78	Select storage pool name	324
7-79	Save Image - Create ISO file	325
7-80	Select the output destination folder	326
7-81	Change file type to .iso	326
7-82	Blade Task - Remote Control	328
7-83	Start Remote Control	328
7-84	Remote Control Java window	328
7-85	Select image option	329
7-86	Browse and select file	329
7-87	File added to Selected Resources list	330
7-88	Selected Resources mounted	330
7-89	New physical optical devices	331

7-90 IBM i power down partition	333
8-1 Remote Control window - assign Media Tray	337
8-2 BladeCenter System status.....	338
8-3 Activating an IVM partition	338
8-4 Opening a terminal window from the IVM.....	339
8-5 SMS menu	340
8-6 Activating an IVM partition	342
8-7 Opening a terminal window from the IVM.....	342
8-8 SMS menu	343
8-9 Red Hat Enterprise Linux 5.3 welcome screen.....	344
8-10 Select network device	345
8-11 TCP/IP configuration panel.....	346
8-12 TCP/IP configuration of IP address, gateway, and name server	346
8-13 NFS server configuration window panel	347
8-14 Select between VNC or text installation modes panel	347
8-15 Installation number panel	348
8-16 Select Partitioning Type panel	349
8-17 Review Partitioning panel	350
8-18 Select additional packages panel	351
8-19 Installation progress window.....	351
8-20 Installation complete panel	352
8-21 Setup Agent panel.....	353
8-22 Finished Red Hat Enterprise Linux 5.3 installation	353
0-1 Kickstart main window with Basic Configuration panel (©2008 Red Hat, Inc.).....	357
0-2 Installation Method panel (©2008 Red Hat, Inc.)	358
0-3 Partition Information panel (©2008 Red Hat, Inc.)	359
0-4 Kickstart Configurator Network Configuration panel (©2008 Red Hat, Inc.) .	360
0-5 Firewall Configuration panel (©2008 Red Hat, Inc.)	361
0-6 Package Selection panel (©2008 Red Hat, Inc.).....	362
8-23 Open Firmware prompt	365
9-1 Start Remote Console panel.....	370
9-2 The checkmark indicates that the Media Tray (MT) is owned by the blade in Bay 1	370
9-3 Activating an IVM partition	371
9-4 Opening a terminal window from the IVM	371
9-5 SMS Menu	372
9-6 Select CD-ROM	372
9-7 Select Mode Boot	373
9-8 Select installation type	373
9-9 Activating an IVM partition	374
9-10 Opening a terminal window from the IVM	375

9-11	SMS Menu	376
9-12	Main Menu	377
9-13	Expert	377
9-14	Load ppc Modules	378
9-15	Start installation or update option	378
9-16	Static network configuration example	379
9-17	YaST Welcome panel	380
9-18	Installation Mode	381
9-19	Clock and Time Zone	381
9-20	Installation Settings	382
9-21	AGFA License Agreement	383
9-22	Confirm Installation	383
9-23	YaST installation progress window	384
9-24	Confirm hardware detection window	385
9-25	Reboot now	385
9-26	root User Password	386
9-27	Hostname and Domain Name	387
9-28	Change network configuration	388
9-29	Select firewall	388
9-30	Services to allow list and selecting SSH service	389
9-31	Allowed Service Secure Shell Server (SSH)	390
9-32	SSH port is open	390
9-33	User Authentication Method	391
9-34	New Local User	392
9-35	Installation completed window	393
9-36	Login screen	394
9-37	Open Firmware prompt	396
10-1	BladeCenter Power Domain Summary	399
10-2	Additional power settings	400
10-3	Chassis thermal and trending options	400
10-4	Power Management Policies	401
10-5	Power Domain Details	402
10-6	Blade power configuration settings	403
10-7	Bladeserver trend data	404
10-8	Director menu options	409
10-9	Active Energy Manager options	410
10-10	Select resource	411
10-11	Actions options	411
10-12	Properties - Active Energy tab	412
10-13	Edit values	412
10-14	Power Capping option	413
10-15	Power Capping options	414
10-16	power capping enabled	414

10-17	Power Savings option	415
10-18	Power Savings options	416
10-19	Trend Data option	416
10-20	Trend Data display	417
10-21	Trend data chart options	418
10-22	Information event details	418
10-23	Export data	419
10-24	Energy calculator option	420
10-25	Energy cost properties	420
10-26	Calculated energy cost	421
10-27	Work with power policies	422
10-28	Select targets - browse	422
10-29	Group Select	423
10-30	Select targets	423
10-31	Selected targets added	424
10-32	Power policy wizard welcome	425
10-33	Policy name and description	425
10-34	Power policy type	426
10-35	Power policy settings	426
10-36	Power policy summary	427
10-37	Available power policies	427
10-38	Apply power policy	428
10-39	Run now - policy apply option	428
10-40	Policy schedule options	429
10-41	Notification tab	430
10-42	Policy options tab	431
11-1	Management Partition Updates view	435
11-2	PowerVM Enterprise key entry	437
11-3	Memory region size	439
11-4	hidsk reserve policy not set correctly	441
11-5	Partition Migration validation error message for target storage	442
11-6	Processor compatibility mode on JS23/JS43	445
11-7	Processor compatibility mode on JS12/JS22	446
11-8	Change the processor compatibility mode on JS23/JS43	447
11-9	Virtual optical device to be removed	448
11-10	Partition Migrate option	449
11-11	Partition Mobility validation	450
11-12	Partition Migration validation success	451
11-13	Failed validation because of improper virtual SCSI device assignment	452
11-14	Migrate Status view	453
11-15	Starting the status view for a mobile partition	454
11-16	Remote IVM indicating migration in progress	455
11-17	Partition migration complete to remote system	456

12-1	Select BladeCenter boot mode main page	463
12-2	Firmware selection page	464
12-3	Blade Power / Restart	465
12-4	Enter SMS Menu	466
12-5	SMS main menu	467
12-6	SMS firmware boot side options	468
12-7	SMS exit page	468
12-8	Firmware Vital Product Data (VPD) on AMM	469
12-9	Firmware level inside the SMS main menu	470
12-10	IBM BladeCenter JS23 and JS43 Control Panel	473
12-11	AMM BladeCenter LEDs control and status page	475
12-12	LEDs on the IBM BladeCenter JS23 Express board. This is the same LED configuration of the IBM BladeCenter JS43 Express Base planar .	476
12-13	LEDs on the IBM BladeCenter JS43 Express MPE planar	477
12-14	AMM Blade Server Data panel	481
12-15	Power-on checkpoints inside AMM web interface	482
12-16	AMM event log interface	483
12-17	Partitions view	484
12-18	More Tasks - Reference Codes	485
12-19	Reference Code list - normal IPL	486
12-20	IBM i partition error	487
12-21	Reference Code list - error condition	488
12-22	AMM Event Log	489
12-23	Event log filter	489
12-24	Event log data details	490
12-25	Monitor tasks	490
A-1	JS23/JS43 Control Panel	495
A-2	AMM login panel	497
A-3	Select time-out parameter	498
A-4	Blade Remote Control options	499
A-5	Start remote control	499
A-6	Remote control - remote console and remote disk	500
A-7	Power/Restart in the left menu of the AMM Web interface	502
A-8	Power/Restart blade options	503
A-9	Serial over a LAN console - select active console	506
A-10	Physical console shown with remote control - select active console .	507
A-11	JS23/JS43 SMS	508
B-1	YaST Control Center in graphics mode	524
B-2	YaST Control Center in text mode	525
B-3	Selecting the Autoinstallation option	526
B-4	Main AutoYaST menu (SLES 11)	527
B-5	Create a reference profile	528
B-6	Selecting additional resources	529

B-7	AutoYaST software selection	530
B-8	Configure the root user	532
C-1	Configure a TFTP server in SLES11	539
C-2	Initial setup of SLES NFS installation server	542
C-3	Source configuration window	543
D-1	Service diagnostic and productivity packages decision tree	547
D-2	Select “on BladeCenter servers”	548
D-3	OS level selection tabs example	549
D-4	Available packages for Red Hat on BladeCenter servers	549
D-5	Available packages for Red Hat/SuSE Linux on IVM-managed server .	552
D-6	DLPAR and Live Partition mobility services are enabled	555

Tables

2-1 Processor features	11
2-2 Memory features	11
2-3 Storage features	12
2-4 Virtualization features	12
2-5 Predictive failure analysis features	12
2-6 Environment considerations	13
2-7 Physical characteristics	13
2-8 Supported I/O options	14
2-9 Integrated Functions	15
2-10 Integrated functions	15
2-11 Supported operating systems	16
2-12 JS23 standard configuration	19
2-13 BladeCenter JS23 processor support	20
2-14 BladeCenter JS23 memory support	20
2-15 BladeCenter JS23 disk support	21
2-16 JS43 configuration options	21
2-17 BladeCenter JS43 processor support	23
2-18 BladeCenter JS43 memory support	23
2-19 BladeCenter JS43 disk support	23
2-20 BladeCenter JS23 and JS43 chassis support	24
2-21 BladeCenter support	25
2-22 BladeCenter H specifications	28
2-23 BladeCenter S specifications	32
2-24 BladeCenter HT specifications	35
2-25 Power Consumption for IBM BladeCenter JS23 and JS43 Express	37
3-1 Supported High Speed Daughter Cards (HSDC)	54
3-2 Supported CIOv PCI-e Expansion Cards	55
3-3 PowerVM editions for IBM BladeCenter JS23 and JS43 Express	62
4-1 Optical drive location paths	133
5-1 Estimated additional VIOS CPU entitlement per shared memory LPAR	182
7-1 BladeCenter JS23 feature list	245
7-2 BladeCenter JS43 feature list	246
12-1 Description of Front Panel buttons and LEDs	474
12-2 Lightpath LED description	477
12-3 Types of SRCs	480

Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing, IBM Corporation, North Castle Drive, Armonk, NY 10504-1785 U.S.A.

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurement may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming

techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs.

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. These and other IBM trademarked terms are marked on their first occurrence in this information with the appropriate symbol (® or ™), indicating US registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the Web at <http://www.ibm.com/legal/copytrade.shtml>

The following terms are trademarks of the International Business Machines Corporation in the United States, other countries, or both:

1350™	i5/OS®	Predictive Failure Analysis®
Active Memory™	IBM Systems Director Active	Redbooks®
Advanced Peer-to-Peer Networking®	Energy Manager™	Redbooks (logo)  ®
AIX 5L™	IBM®	RETAIN®
AIX®	iSeries®	ServeRAID™
BladeCenter®	Micro-Partitioning™	ServerProven®
Chipkill™	OS/400®	System i™
DS4000™	POWER™	System p™
DS6000™	POWER Hypervisor™	System p5™
DS8000™	POWER5™	System Storage™
EnergyScale™	POWER6™	Tivoli®
eServer™	POWER6+™	Wake on LAN®
	PowerVM™	

The following terms are trademarks of other companies:

Advanced Micro Devices, AMD, ATI, the AMD Arrow logo, and combinations thereof, are trademarks of Advanced Micro Devices, Inc.

InfiniBand, and the InfiniBand design marks are trademarks and/or service marks of the InfiniBand Trade Association.

ITIL is a registered trademark, and a registered community trademark of the Office of Government Commerce, and is registered in the U.S. Patent and Trademark Office.

Novell, SUSE, the Novell logo, and the N logo are registered trademarks of Novell, Inc. in the United States and other countries.

QLogic, and the QLogic logo are registered trademarks of QLogic Corporation. SANblade is a registered trademark in the United States.

ABAP, SAP NetWeaver, SAP, and SAP logos are trademarks or registered trademarks of SAP AG in Germany and in several other countries.

Java, JRE, Power Management, Solaris, Sun, and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

Excel, Microsoft, Windows Server, Windows Vista, Windows, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

Intel, Intel logo, Intel Inside logo, and Intel Centrino logo are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Linux is a trademark of Linus Torvalds in the United States, other countries, or both.

Other company, product, or service names may be trademarks or service marks of others.

Preface

This IBM® Redbooks® publication provides a detailed technical guide for configuring and using the IBM BladeCenter® JS23 and IBM BladeCenter JS43 servers. These IBM Power Blade servers feature the latest IBM POWER6™ processor technology. This book teaches you how to set up the latest Power Blade servers to run AIX®, i, and Linux® operating systems in the IBM BladeCenter architecture.

This book will be useful for data center system architects; network, storage and facilities engineers; system administrators; and application architects that need a detailed system bring-up guide and integration guide for IBM Power Blade servers hosting virtualized operating system partitions. We include procedures for installation and configuration of Virtual Input/Output Server (VIOS), Integrated Virtualization Manager (IVM), IBM AIX Version 6.1, IBM i 6.1 (i5/OS® V6R1), Red Hat Enterprise Linux, and SUSE® Enterprise Linux.

The team that wrote this book

This book was produced by a team of specialists from around the world working at the International Technical Support Organization, Austin Center.

Alex Zanetti de Lima is a Software Engineer at the Linux Technology Center in Brazil since 2007 and leads the development of the IBM Installation Toolkit project. He has experience on developing, debugging and deploying SLES and RHEL Linux on Bladecenters JS23 and JS43, He's also experienced on developing and deploying server consolidation solutions using Bladecenters and holds two patents in the United States Patent and Trademark Office related to this area.

Kerry Anders is a Consultant in System p™ Lab Services for the IBM Systems and Technology Group, based in Austin, Texas. He supports clients implementing POWER™-based processor blades using Virtual I/O Server, Integrated Virtualization Manager, and AIX. He was previously the Systems Integration Test team lead for the IBM BladeCenter JS21blade with IBM SAN storage using AIX and Linux. His prior work includes test experience with the JS20 blade also using AIX and Linux in SAN environments. Kerry began his career with IBM in the Federal Systems Division supporting NASA at the Johnson Space Center as a systems engineer. He transferred to Austin in 1993.

Nahman Cohen is an IT manager in Memorex Telex Israel for 10 years. He has 18 years of experience in the network support and hardware fields. His areas of expertise include Windows®, Linux, Sun™ Solaris™ and networks.

Steve Strain is a Software Engineer/Advisory Education Specialist for the Rochester Support Center in Rochester MN. He is responsible for developing and delivering education based on IBM i and the POWER platforms. He has worked on many HMC, POWER5™ and POWER6 courses tailored towards hardware maintenance, integration and system configuration. Prior to his education position he was a member of the Rochester Support Center hardware team specializing in DASD systems diagnostics and repair. He has been working for IBM since 1992.

Vasfi Gucer is an IBM Certified Consultant IT Specialist at the ITSO Austin Center. He has more than 15 years of experience on teaching and implementing systems management, networking hardware, and distributed platform software. Vasfi is also an ITIL® Expert and PMP.

Acknowledgements

This IBM Redbooks publication would not have been possible without the generous support of many IBM employees around the world. For their contributions, guidance, patience, and technical feedback in support of this project, we gratefully acknowledge the following IBM employees:

Chris Austen, John E Biebelhausen, David Engebretsen, Kaena Freitas, Anita Govindjee, Carol Hernandez, Hoa Nguyen, Vess Natchev, Jonathan Van Nieuwaal, Gary Rohret, Todd Rosedahl, Steven Royer, Mark Smolen, Mike Stys, Padmavathy Tamirisa, Mark Vanderwiel, Mary Wigham, Kyle Wurgler.

We would also like to thank the authors of these previous IBM Redbooks publications: *IBM BladeCenter JS21: The POWER of Blade Innovation*, SG24-7273, *IBM BladeCenter JS21 Technical Overview and Introduction*, REDP-4130, *Techniques to Build a Diskless Boot Linux Cluster of JS21 Blades*, REDP-4238, *IBM BladeCenter JS12 and JS22 Implementation Guide*, SG247655, *IBM PowerVM Live Partition Mobility*, SG247460

Chris Almond, Kerry Anders, John E Bailey, Trina Bunting, Benjamin Ebrahimi, Hendrik Ernst, Stephen Hochstetler, Tom Junkin, Wayne Kimble, Gregor Linzmeier, Bob Maher, Ameet Paranjape, Thomas Prokop, Guido Somers, Scott Vetter and Bernhard Zeller.

Lab systems setup in support of this project required the availability of multiple IBM BladeCenter chassis, POWER blade servers, plus various adapters and access to disk storage subsystems. We are very thankful for the lab systems support we received from:

Ned Gamble and **Erich J Hauptli**.

Finally, the team would also like to acknowledge the support for this project provided by **Scott Vetter**, ITSO System p Team Leader; and our book editor, **Wade Wallace**, also contributed to our production and review efforts.

Become a published author

Join us for a two- to six-week residency program! Help write a book dealing with specific products or solutions, while getting hands-on experience with leading-edge technologies. You will have the opportunity to team with IBM technical professionals, Business Partners, and Clients.

Your efforts will help increase product acceptance and customer satisfaction. As a bonus, you will develop a network of contacts in IBM development labs, and increase your productivity and marketability.

Find out more about the residency program, browse the residency index, and apply online at:

ibm.com/redbooks/residencies.html

Comments welcome

Your comments are important to us!

We want our books to be as helpful as possible. Send us your comments about this book or other IBM Redbooks in one of the following ways:

- ▶ Use the online **Contact us** review Redbooks form found at:

ibm.com/redbooks

- ▶ Send your comments in an e-mail to:

redbooks@us.ibm.com

- ▶ Mail your comments to:

IBM Corporation, International Technical Support Organization
Dept. HYTD Mail Station P099

2455 South Road
Poughkeepsie, NY 12601-5400



Part 1

The IBM BladeCenter JS23 and JS43 servers

This Part provides general and technical descriptions of the BladeCenter products covered in this publication.



Introduction to IBM BladeCenter

This chapter provides an introduction to IBM BladeCenter and blade servers JS23 and JS43 and discuss the business benefits of blade servers in general and has the following sections:

- ▶ “Highlights of BladeCenter” on page 4
- ▶ “IBM BladeCenter is the right choice, open, easy and green” on page 6

1.1 Highlights of BladeCenter

Blade servers are thin servers that insert into a single rack-mounted chassis which supplies shared power, cooling, and networking infrastructure. Each server is an independent server with its own processors, memory, storage, network controllers, operating system, and applications. Blade server design is optimized to minimize physical space. Standard rack-mount servers require a number of power cords and network cables, but in the case of blade servers, blade enclosures provide these services such as power, cooling, networking for multiple blade servers, thereby reducing the space and cable requirements. See Figure 1-1 on page 4 for a computing environment with and without blade servers.

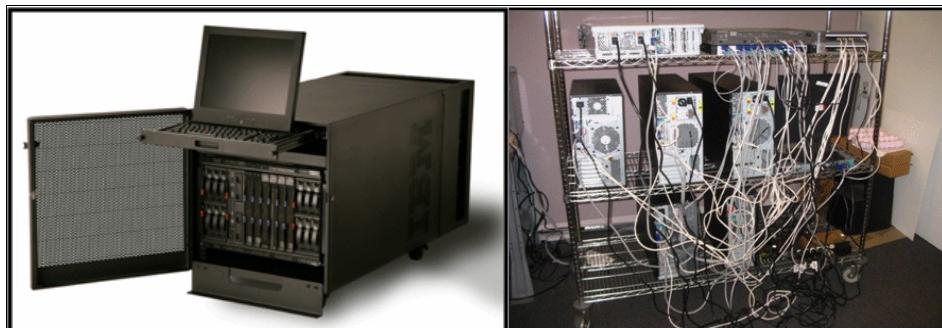


Figure 1-1 With and without blade servers

Blade servers came to market around 2000, initially to meet clients' needs for greater ease of administration and increased server density in the data center environment. When IBM released the IBM BladeCenter in November 2002, it quickly changed the industry with its modular design. The IBM BladeCenter provides complete redundancy in a chassis, enables network and storage integration. IBM BladeCenter has the following strong is ideal for infrastructure consolidation, virtualization, and demanding applications which require scalable performance and high memory capacity. The IBM BladeCenter blade server processors use POWER6 processor technology and has the ability to run AIX, IBM System i™, and Linux operating systems simultaneously. Overall IBM BladeCenter provides a secure, resilient and dynamic infrastructure solution that helps drive cost down, reduces risk, improves energy efficiency and enhances flexibility.

Built on the promise of the IBM BladeCenter family of products-easy-to-use, integrated platforms with a high degree of deployment flexibility, energy efficiency, scalability and manageability-the BladeCenter JS23 and JS43 Express are the premier blades for 64-bit applications. They represent one of the

most flexible and cost-efficient solutions for UNIX®, i and Linux deployments available in the market. Further enhanced by its ability to be installed in the same chassis with other IBM BladeCenter blade servers, the JS23 and JS43 can deliver the rapid return on investment that clients and businesses demand.

Delivering on the promise of a truly dynamic infrastructure the BladeCenter JS23 and JS43 help in delivering superior business and IT services with agility and speed - all in a simple to manage highly efficient way.

The JS23 and JS43 Express blades have been pre-configured and tested by IBM and are based on proven technology. Utilizing a 4.2 GHz 64-bit POWER6 processor and available in a 4-core or 8-core configuration, they are designed to deliver outstanding performance and capabilities at compelling prices. With faster and more reliable double data rate 2 (DDR2) memory options and support for eight to sixteen memory dual in-line memory module (DIMM) slots along with Serial Attached SCSI (SAS) disk subsystem, the BladeCenter JS23 and JS43 Express blades are designed for increased uptime and enhanced performance. Exploiting the newest in high performance and energy efficient solutions like Solid State Disk Technology which offers remarkable performance for I/O intensive applications and those that require fast and frequent data access the JS23 and JS43 offers the utmost in easy to manage, tough to break solutions. Along with built-in support for PowerVM™ Editions for advanced virtualization, the JS23 and JS43 offers an ideal blade server solution for driving your most demanding performance and memory-intensive workloads such as virtualization, databases and high performance computing (HPC) applications. By consolidating and virtualizing on BladeCenter with the JS23 and JS43, you can increase the utilization of your hardware and decrease the number of physical assets you need to watch over. This translates into real dollar savings through better energy conservation and IT resource usage across the data center.

Meeting today's cost challenges while planning for tomorrow's growth can be difficult. When you demand more performance, this blade server is ready to scale to four processors with the addition of a multi-processor expansion unit (MPE). Simply snap the MPE to the original blade and it becomes a 4-processor 8 core, ~60 mm blade server with the additional performance, memory and double the I/O expansion you need to deliver a flexible, resilient and highly scalable IT infrastructure. This expansion capability is highly cost effective and provides the opportunity to standardize on a single blade platform for all your needs.

With IBM BladeCenter and System i, clients can realize innovation with proven 6th generation POWER technology for enhanced performance, efficiency and reliability; cut costs and consolidate workloads with leadership virtualization as a built-in feature; go green with better energy management by using EnergyScale™ technology; and manage growth, complexity and risk with the flexibility of IBM BladeCenter technology. With System i support, small and medium clients can consolidate their System i and Intel® processor-based

servers into a single chassis, leveraging the management, space and power savings provided by IBM BladeCenter solutions. Large or small enterprises can now consolidate their older i5/OS applications into a centralized BladeCenter environment with a choice of BladeCenter chassis and blade configurations to fit their needs.

Simplify, cut costs, boost productivity, go green. These are all priorities for IT, and they are all driving organizations to rethink their server strategies and become more receptive to new ways to use IT. Blades are the next-generation solution, promising improvements across the board. The IBM BladeCenter innovative, open design offers a true alternative to today's sprawling racks and overheated server rooms. Migrating to the blade solution will give you a solution that uses less energy and more choices and control with less complexity.

1.2 IBM BladeCenter is the right choice, open, easy and green

IBM BladeCenter is the right choice, open, easy and green.

The RIGHT choice, tailored to fit your diverse needs:

- ▶ It is flexible and modular. As needs evolve, a one-size-fits-all solution doesn't work.
 - Meet your needs with BladeCenter: everything from a high-performance data center to a small office with limited IT skills-IBM has you covered
 - Get flexibility with 5 compatible chassis and 5 blade types supporting multiple I/O fabrics, all managed from a common point
- ▶ It is robust and reliable, providing redundancy throughout and the information you need to keep your business up and running.
 - Provide redundancy for no single point of failure with IBM BladeCenter
 - Preserve application uptime with IBM Predictive Failure Analysis® and light path diagnostics
 - Make decisions based on accurate data for quick problem diagnosis with First Failure Data Capture (FFDC)

OPEN and innovative, for a flexible business foundation:

- ▶ It is comprehensive, providing broad, fast, and reliable networking and storage I/O with BladeCenter Open Fabric.
 - Match your data center needs and the appropriate interconnect using a common management point, and 5 I/O fabrics to choose from

- Extract the most from your third-party management solutions by utilizing the BladeCenter Open Fabric Manager
- ▶ It is collaborative, enabling you to harness the power of the industry to deliver innovation that matters.
 - Get flexibility from a myriad of solutions created by Blade.org members and industry leaders that have downloaded our open specification

EASY to deploy, integrate and manage:

- ▶ It enables efficient integrated management, which allows you to minimize costs with the tools you need for effective management.
 - Automate OS installation and BIOS updates remotely with IBM Director tools
 - Administer your blades at the chassis or rack level with the Advanced Management Module
 - Plug into your enterprise management software
- ▶ It enables deployment simplicity without trade-offs by speeding the deployment of new hardware in minutes rather than days, using BladeCenter Open Fabric Manager
 - Get significantly faster deployment of servers and I/O than from rack solutions
 - Reduce costly downtime with integrated failover capability
 - Manage from a single point of control via the Advanced Management Module
 - Use with virtually all IBM switches, blades and chassis

GREEN today for a better tomorrow:

- ▶ It offers control via powerful tools that help you optimize your data center infrastructure so you can be responsive.
 - Understand your power requirements with IBM Power Configurator
 - Monitor, control and virtualize your power with IBM Systems Director Active Energy Manager™ for x86
 - Reduce data center hot spots with the IBM Rear Door Heat eXchanger
 - Optimize and future-proof your data center with IBM Data Center Energy Efficiency services
- ▶ Our eco-friendly servers and services can help you be environmentally responsible.
 - Become more energy efficient with IBM expertise



General description

The newest release of the IBM BladeCenter POWER6 processor based blade family consists of two new models: The JS23 and JS43 Express blade servers.

This chapter provides an overview of these 2 new blade servers and has the following sections.

- ▶ “Overview of the JS23 and JS43 Express blade servers” on page 10
- ▶ “Features and specifications” on page 10
- ▶ “Blade Server features” on page 16
- ▶ “Physical specifications BladeCenter JS23” on page 18
- ▶ “Physical specifications BladeCenter JS43” on page 21
- ▶ “IBM BladeCenter chassis” on page 24

2.1 Overview of the JS23 and JS43 Express blade servers

The newest release of the IBM BladeCenter POWER6 processor based blade family consists of two new models: The JS23 and JS43 Express blade servers. This new family allows processor scalability starting with a 2 processor (4-core single wide) blade and adds the ability to upgrade to a 4 processor (8-core) blade with the addition of a second blade making it a double wide package.

The new blades continue to support AIX, IBM i, and Linux operating systems. Also supported are the IBM EnergyScale technologies, IBM PowerVM Virtualization and AltiVec SIMD acceleration functionality.

The optional IBM Systems Director Active Energy Manager (AEM) for POWER V5.1 and browser-based Integrated Virtualization Manager (IVM) software make it easier than ever to achieve increased utilization and energy efficiency using the PowerVM and EnergyScale capabilities of the system.

The IBM BladeCenter JS23/JS43 Express blade servers refresh every major subsystem to the latest industry-standard functionality. With the energy-efficient 4.2 GHz high performance dual-core, 64-bit POWER6 processors along with the improved reliability of a SAS disk drive bay supporting multiple disk sizes, the JS23/JS43 combine leadership performance per watt and enterprise-class reliability features. The BladeCenter JS23 Express supports up to 64 GB of DDR2-based memory in eight DIMM slots. The BladeCenter JS43 Express supports up to 128 GB of DDR2-based memory in sixteen DIMM slots. In addition to standard support for ECC and IBM Chipkill™ technology designed for improved memory fault protection, the BladeCenter JS23/JS43 Express also offers memory running at up to 667 MHz for outstanding performance.

The BladeCenter JS23 Express (7778-23X), BladeCenter JS43 Express (7778-23X + FC8446), BladeCenter S Chassis (7779-BCS), and BladeCenter H Chassis (7989-BCH) are part of the Cluster 1350™ portfolio.

The JS23 and JS43 blades can be configured in the IBM eConfigurator similar to other IBM System p servers.

2.2 Features and specifications

The following tables provide some of the features and specifications of the IBM BladeCenter JS23 and JS43 blades.

Table 2-1 on page 11 provides a general overview of the processor features of the IBM BladeCenter JS23 and JS43.

Table 2-1 Processor features

Component	Features
Microprocessor	JS23: Two dual-core (4-way) 64-bit POWER6 microprocessors; 4.2 GHz JS43: Two additional dual-core (total 8-way) 64-bit POWER6 microprocessors; 4.2 GHz Support for Energy Scale thermal management for power management/oversubscription (throttling) and environmental sensing

Table 2-2 on page 11 provides a general overview of the memory features of the IBM BladeCenter JS23 and JS43.

Table 2-2 Memory features

Component	Features
Memory	JS23: Dual-channel (DDR2) with 8 very low profile (VLP) DIMM slots (maximum 64 GB) JS43: Dual-channel (DDR2) with 16 VLP DIMM slots (maximum 128 GB) Supports 2 Gigabyte (GB), 4 GB, and 8 GB DDR2 DIMMs Supports 2-way interleaved, DDR2, PC2-4200 or PC2-5300, ECC SDRAM registered x4, memory scrubbing, Chipkill, and bit steering DIMMs

Table 2-3 on page 12 provides a general overview of the storage features of the IBM BladeCenter JS23 and JS43.

Table 2-3 Storage features

Component	Features
Storage	JS23: Support for one internal small-form-factor (SFF) Serial Attached SCSI (SAS) drive or Solid State Drive (SSD) in the base unit JS43: Support for one additional internal SFF SAS drive or SSD in the expansion unit for a total of two drives

Table 2-4 on page 12 provides a general overview of the virtualization features of the IBM BladeCenter JS23 and JS43.

Table 2-4 Virtualization features

Component	Features
Virtualization	PowerVM Standard Edition hardware feature supports Integrated Virtualization Manager (IVM) and Virtual I/O Server No HMC support Partition Migration No CUoD

Table 2-5 on page 12 provides a general overview of the predictive failure analysis features of the IBM BladeCenter JS23 and JS43.

Table 2-5 Predictive failure analysis features

Component	Features
Predictive Failure Analysis performed on:	Microprocessor Memory

Table 2-6 on page 13 provides a general overview of the environmental considerations of the IBM BladeCenter JS23 and JS43.

Table 2-6 Environment considerations

Component	Features
Environment	<p>Electrical Input: 12V dc</p> <p>Air temperature: Blade server on: 10° to 35°C (50° to 95°F). Altitude: 0 to 914 m (3000 ft)</p> <p>Blade server on: 10° to 32°C (50° to 90°F). Altitude: 914 m to 2133 m (3000 ft to 7000 ft)</p> <p>Blade server off: -40° to 60°C (-40° to 140°F)</p> <p>Humidity: Blade server on: 8% to 80% Blade server off: 8% to 80%</p>

Table 2-7 on page 13 provides a general overview of the physical characteristics of the IBM BladeCenter JS23 and JS43.

Table 2-7 Physical characteristics

Component	Features
Physical characteristics	<p>JS23 Size: Height: 24.5 cm (9.7 inches)</p> <p>Depth: 44.6 cm (17.6 inches)</p> <p>Width: 30 mm (1.14 inches) single-slot blade</p> <p>Maximum weight: 5.0 kg (11 lb)</p> <p>JS43 Size: Height: 24.5 cm (9.7 inches)</p> <p>Depth: 44.6 cm (17.6 inches)</p> <p>Width: 60 mm (2.28 inches) double-slot blade</p> <p>Maximum weight: 10.0 kg (22 lb)</p>

Table 2-8 on page 14 provides information on supported I/O options for the IBM BladeCenter JS23 and JS43.

Table 2-8 Supported I/O options

Component	Features
I/O adapter card options	Up to two PCIe High Speed adapters on JS43. Only one supported on JS23 Up to two PCIe CIOv adapters on JS43. Only one on JS23

Table 2-9 on page 15 and Table 2-10 on page 15 provide a general overview of the integrated functions of the IBM BladeCenter JS23 and JS43.

Table 2-9 Integrated Functions

Component	Features
Integrated functions	<p>JS23: Two 1 GB Ethernet controllers connected to the BladeCenter chassis fabric through the 5-port integrated Ethernet switch</p> <p>JS43: Two additional 1 GB Ethernet controllers, connecting directly to BladeCenter Ethernet switch modules</p> <p>Expansion card interface</p> <p>The baseboard management controller (BMC) is a flexible service processor with Intelligent Platform Management Interface (IPMI) firmware and SOL support</p> <p>PCI attached ATI™ RN 50 graphics controller</p> <p>SAS controller</p> <p>Light path diagnostics RS-485 interface for communication with the management module</p> <p>Automatic server restart (ASR)</p> <p>Serial over LAN (SOL)</p> <p>Four Universal Serial Bus (USB) buses on base planar for communication with keyboard and removable-media drives</p>

Table 2-10 Integrated functions

Component	Features
Integrated functions continued	<p>Support for USB-attached local keyboard, video, and mouse (KVM)</p> <p>Transferable Anchor function (Renesas Technology HD651330 microcontroller) in the management card</p>

Table 2-11 on page 16 provides information on supported operating systems for the IBM BladeCenter JS23 and JS43.

Table 2-11 Supported operating systems

Component	Features
Operating system	Linux SLES10 SP2 or later versions Red Hat RHEL 5.2 or later versions Red Hat RHEL 4.6 or later versions AIX 5.3.S, 6.1.F IBM i 6.1

2.3 Blade Server features

The following are the main features of JS23 and JS43 blade servers:

- ▶ Baseboard management controller (BMC)

The enhanced BMC for the JS23 blade server and JS43 blade server is a flexible service processor that provides support for the following functions:

- Alert Standard Format (ASF) and RS-485 interfaces to the management modules
 - Intelligent Platform Management Interface (IPMI)
 - The operating system
 - Power control and advanced power management
 - Reliability, availability, and serviceability (RAS) features
 - Serial over LAN (SoL)
 - Continuous health monitoring and control
 - Configurable notification and alerts
 - Event logs that are time stamped and saved in nonvolatile memory and that can be attached to e-mail alerts
 - Point-to-point protocol (PPP) support
 - Remote power control
 - Remote firmware update and access to critical blade server settings
- ▶ Disk drive support

- The JS23 blade server supports one 2.5 inch hard disk drive. The JS43 blade server can support up to two 2.5 inch hard disk drives. The disk drives can be either the small-form-factor (SFF) Serial Attached SCSI (SAS) or the Solid state drive (SSD).
- ▶ IBM Director
 - IBM Director is a workgroup-hardware-management tool that you can use to centrally manage the JS23 blade server and JS43 blade server, including updating the JS23 and JS43 firmware.
- ▶ Integrated network support

The blade server has two integrated (onboard) 1 gigabit-per-second (Gbps) Host Ethernet Adapter (HEA) controllers that provide advanced acceleration features, such as checksum offload, TCP large send, and jumbo frames. Capabilities include virtualized adapter sharing among logical partitions that does not require the shared Ethernet adapters of Virtual I/O Server. TCP advanced features include hardware de-multiplexing and per connection queues.
- ▶ I/O expansion

The blade server has connections on the system board for two optional PCIe expansion cards, such as Fibre Channel and InfiniBand® expansion cards, for adding more network communication capabilities to the blade server.
- ▶ Large system memory capacity

The memory bus in the JS23 blade server supports up to 64 GB of system memory. The memory bus in the JS43 blade server supports up to 128 GB of system memory.
- ▶ Light path diagnostics

Light path diagnostics provides light-emitting diodes (LEDs) to help you diagnose problems. An LED on the blade server control panel is lit if an unusual condition or a problem occurs. If this happens, you can look at the LEDs on the system board to locate the source of the problem.
- ▶ Power throttling

If your BladeCenter unit supports power management, the power consumption of the blade server can be dynamically managed through the management module.
- ▶ Reliability, availability and serviceability (RAS)

The blade server incorporates RAS features that provide capability to monitor, diagnose and correct errors on some of the components within the blade server. Some of the features are:

- Automatic service processor reset and reload recovery for service processor errors
- Automatic server recovery and restart that provides automatic reboot after boot hangs or detection of checkstop conditions
- Automatic server restart (ASR)
- Built-in monitoring for temperature, voltage, hard disk drives, and flash drives
- Checkstop analysis
- Customer-upgradeable basic input/output system (BIOS) code (firmware code)
- Degraded boot support (memory and microprocessors)
- Extended Error Handling (EEH) for adapter failures
- Emergency power off (EPOW) for the blade server and expansion cards
- Environmental monitors and alerts
- First Failure Data Capture (FFDC) for determining failure root cause
- Service processor communication with the management module to enable remote blade server management
- Light emitting diodes (LEDs) for identifying failing customer replaceable units
- Light path diagnostics
- POWER6 Hypervisor (PHYP) partition recovery or partition termination when unrecoverable errors occur
- Power-on self-test (POST)
- Vital product data (VPD) unique identifiers on blade server and all major electronic components with information stored in nonvolatile memory for remote viewing

2.4 Physical specifications BladeCenter JS23

The next sections will provide more details related to the features of the JS23 and JS43 blade servers.

2.4.1 Minimal and optional features of BladeCenter JS23

The BladeCenter JS23 blade is based on a modular design of two dual-core-modules (DCM) that contain 64-bit POWER6 processors and

integrated L2 cache soldered directly to the system planar board. Additionally there is a 32MB L3 cache that is integrated into each of the DCM modules. The JS23 is contained in a single wide package.

Table 2-12 shows the JS23 configuration options.

Table 2-12 JS23 standard configuration

7778-23X	Processor	L2/L3	Memory	Ethernet	Disk
#7778-23X	2-socket, 4-core, 4.2 GHz, POWER6	L2 - 4 MB per core L3 32MB per DCM	4 GB - 64 GB	Dual Gigabit	73,146 or 300GB SAS or 69GB SSD

Figure 2-1 on page 19 shows the physical layout of the JS23 blade including memory slots, disk, and the expansion option connectors.

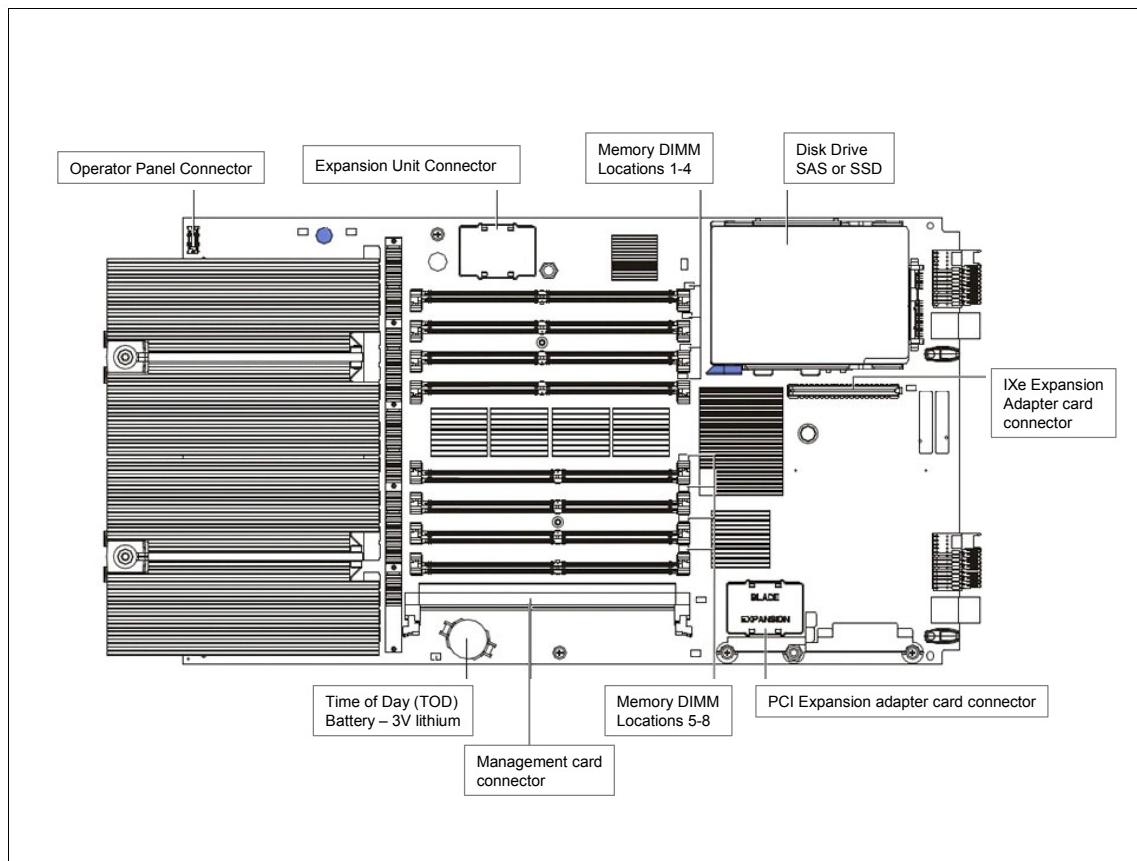


Figure 2-1 JS23 Blade physical layout

2.4.2 Processor features

The key processor features are as follows:

- ▶ The BladeCenter JS23 blade provides the support for a 2-socket, 4-core, POWER6 4.2 GHz processor implementation.
- ▶ Each processor is directly mounted to the system planar board, providing multi-processing capability.
- ▶ Each processor core includes a 64-KB Instruction-Cache, 64-KB Data-Cache, and 4 MB of L2 cache.
- ▶ Each DCM contains a 32MB L3 Cache.

Table 2-13 shows the supported processor on a BladeCenter JS23 blade.

Table 2-13 BladeCenter JS23 processor support

Feature	Description
#7778-23X	IBM BladeCenter JS23 4-core 64 bit 4.2 GHz

2.4.3 Memory features

The integrated memory controller supports eight pluggable registered DIMMs, which must be installed in pairs. The minimum memory that can be installed is 4 GB (2x2 GB) and the maximum is 64 GB (4x16 GB). All the memory features support memory scrubbing, error correction, chipkill, and bit steering. Memory is packaged in Very Low Profile (VLP) RDIMM packages. This newer style memory is much lower in height allowing more memory in the same space.

2.4.4 Memory DIMMs

Table 2-14 provides a list of supported memory on a BladeCenter JS23 blade.

Table 2-14 BladeCenter JS23 memory support

Feature	Description
#8229	4 GB (2x2 GB) DDR2 667 MHz VLP RDIMMs
#8239	8 GB (2x4 GB) DDR2 667 MHz VLP RDIMMs
#8245	16 GB (2x8 GB) DDR2 400 MHz VLP RDIMMs

2.4.5 Internal disk

Table 2-15 provides a list of supported disks on a BladeCenter JS23 blade. Disk drives are not required on the base offering.

Table 2-15 *BladeCenter JS23 disk support*

Feature	Description
#8237	73 GB SAS 10K SFF hard disk drive
#8236	146 GB SAS 10K SFF hard disk drive
#8274	300 GB SAS 10K SFF hard disk drive
#8273	69 GB Solid State Disk (SSD)

2.5 Physical specifications BladeCenter JS43

In this section we discuss the physical specifications BladeCenter JS43.

2.5.1 Minimal and optional features of BladeCenter JS43

The BladeCenter JS43 blade is based on a modular design of four dual-core-modules (DCM) that contain 64-bit POWER6 processors and integrated L2 cache soldered directly to the system planar board. Additionally there is a 32MB L3 cache is integrated into each of the DCM modules. The JS43 is contained in a double wide package. The JS43 can be obtained by either starting with a JS23 (7778-23x) and adding a Feature 8446 which contains an additional blade to make a double wide package or ordering a Feature 7778-43X which contains the JS23 (7778-23X) and Feature 8446 assembled at the factory.

Table 2-16 shows the JS43 configuration options.

Table 2-16 *JS43 configuration options*

7778-23X	Processor	L2/L3	Memory	Ethernet	Disk
#7778-23X Plus #8446 Add on expansion module	4-socket, 8-core, 4.2 GHz, POWER6	L2 - 4 MB per core L3 32MB per DCM	8GB - 128 GB	Dual Gigabit	73,146 or 300GB SAS or 69GB SSD

Figure 2-2 on page 22 shows the physical layout of the JS43 blade Multiple Expansion Unit (MPE) including memory slots, disk, and the expansion option connectors. The MPE stacks on top of the single wide JS23 making a double wide blade. Each section has its own processors, memory, disk, and adapter cards.

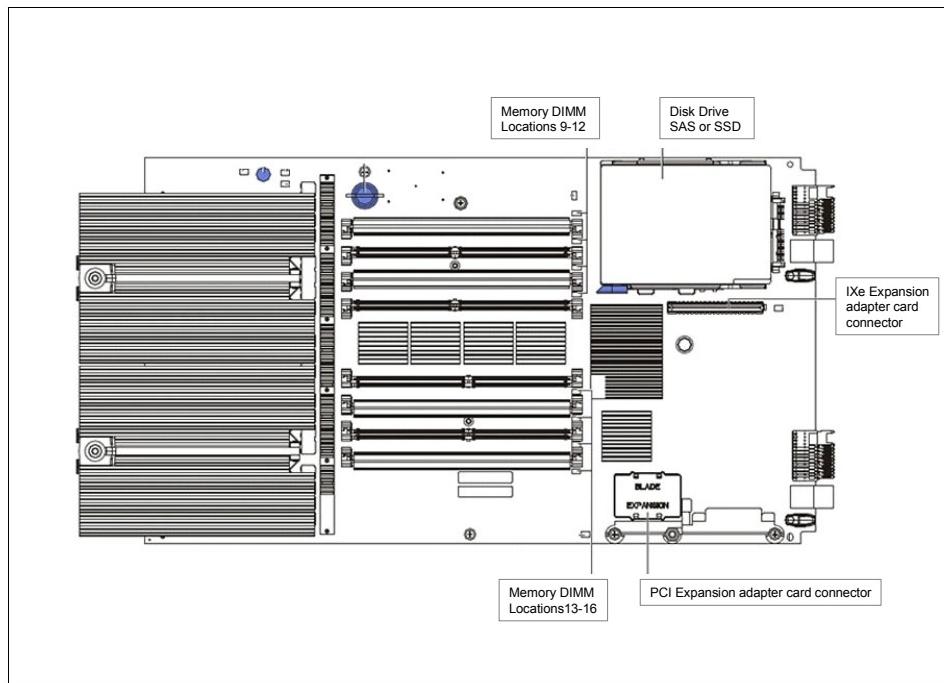


Figure 2-2 JS43 Multiple Expansion Unit (MPE)

2.5.2 Processor features

The key processor features are as follows:

- ▶ The BladeCenter JS43 blade provides the support for a 2-socket, 4-core, POWER6 4.2 GHz processor implementation. Each of the two blade boards contain processors.
- ▶ Each processor is directly mounted to the system planar board, providing multi-processing capability.
- ▶ Each processor core includes a 64-KB Instruction-Cache, 64-KB Data-Cache, and 4 MB of L2 cache.

Each DCM contains a 32MB L3 Cache. Table 2-17 shows the supported processor on a BladeCenter JS43 blade.

Table 2-17 BladeCenter JS43 processor support

Feature	Description
#7778-23X Plus 8446	IBM BladeCenter JS43 8-core 64 bit 4.2 GHz

2.5.3 Memory features

The integrated memory controller supports sixteen pluggable registered DIMMs, which must be installed in pairs. The minimum memory that can be installed is 4 GB (2x2 GB) and the maximum is 128 GB (8x16 GB). All the memory features support memory scrubbing, error correction, chipkill, and bit steering.

2.5.4 Memory DIMMs

Table 2-18 provides a list of supported memory on a BladeCenter JS43 blade.

Table 2-18 BladeCenter JS43 memory support

Feature	Description
#8229	4 GB (2x2 GB) DDR2 667 MHz VLP RDIMMs
#8239	8 GB (2x4 GB) DDR2 667 MHz VLP RDIMMs
#8245	16 GB (2x8 GB) DDR2 400 MHz VLP RDIMMs

2.5.5 Internal disk

Table 2-19 provides a list of supported disks on a BladeCenter JS43 blade. Disk drives are not required on the base offering.

Table 2-19 BladeCenter JS43 disk support

Feature	Description
#8237	73 GB SAS 10K SFF hard disk drive
#8236	146 GB SAS 10K SFF hard disk drive
#8274	300 GB SAS 10K SFF hard disk drive
#8273	69 GB Solid State Disk (SSD)

When you have a JS43 and both halves contain disk drives, RAID is supported. You can utilize either RAID 0 (striping) or RAID 1 (mirroring). Either the SAS

drives or the SSD disk units can be RAIDed however, the drives must be of the same type. It is also preferred to have drives of the same capacity but, RAID can be performed using dissimilar capacities. If differing capacities are used you will only have the effective capacity of the smaller drive.

2.6 IBM BladeCenter chassis

The BladeCenter JS23 and BladeCenter JS43 Express blade are supported in the BladeCenter chassis as shown in Table 2-20. Note that operating system selection (in particular IBM i) should be taken into account for overall system support.

Table 2-20 BladeCenter JS23 and JS43 chassis support

Blade	BC H	BC S	BC HT	BC T	BC E
JS23	YES	YES	YES	NO	NO
JS43	YES	YES	YES	NO	NO

Three chassis are available in the BladeCenter chassis family:

- ▶ IBM BladeCenter H delivers high performance, extreme reliability, and ultimate flexibility for the most demanding IT environments.
- ▶ IBM BladeCenter S combines the power of blade servers with integrated storage, all in an easy-to-use package designed specifically for the office and distributed enterprise environment.
- ▶ IBM BladeCenter HT models are designed for high-performance flexible telecommunication environments by supporting high-speed inter-networking technologies such as 10 Gb Ethernet.

Table 2-21 provides a list of supported BladeCenter chassis and the total number of JSXX blades installable into a chassis. This table describes the physical limitations of blades in chassis.

Note: There are limitations to the number of blade servers that can be installed into the chassis. It is dependent on the power supply configuration, power supply input (110/220V) and power reduction/redundancy options. See “Number of IBM BladeCenter JS23 and JS43 Express in Supported Blade Center Chassis” on page 37 for more information.

Table 2-21 BladeCenter support

Chassis	Number of JS23 blades	Number of JS43 Blades
BladeCenter S chassis	6	3
BladeCenter H chassis	14	7
BladeCenter HT chassis	12	6

2.6.1 BladeCenter H

IBM BladeCenter H delivers high performance, extreme reliability, and ultimate flexibility to even the most demanding IT environments. In a 9U rack space, the BladeCenter H chassis can contain up to 14 blade servers, 10 I/O modules, and four power supplies to provide the necessary I/O network switching, power, cooling, and control panel information to support the individual servers.

The chassis supports up to four traditional fabrics using networking switches, storage switches, or pass through devices. The chassis also supports up to four high-speed fabrics for support of protocols such as 4X InfiniBand or 10 Gigabit Ethernet. The built-in media tray includes light path diagnostics, two front USB inputs, and a DVD drive.

Figure 2-3 and Figure 2-4 on page 27 display the front and rear view of an IBM BladeCenter H.

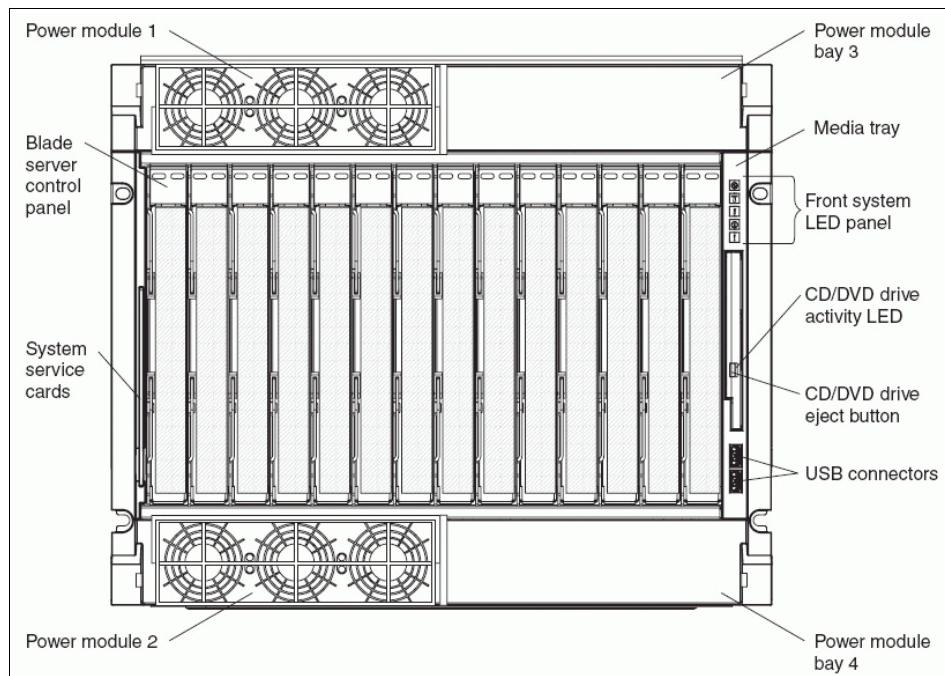


Figure 2-3 Front view of BladeCenter H

The key features on the front of the BladeCenter H are:

- ▶ A media tray at the front right, with a DVD drive, two USB v2.0 ports, and a system status LED panel.
- ▶ One pair of 2,900-watt power modules. An additional power module option (containing two 2,900 W power modules) is available.
- ▶ Two hot swap fan modules (two extra hot swap fan modules are included with the additional power module option).
- ▶ 14 hot swap blade server bays supporting different blade server types.

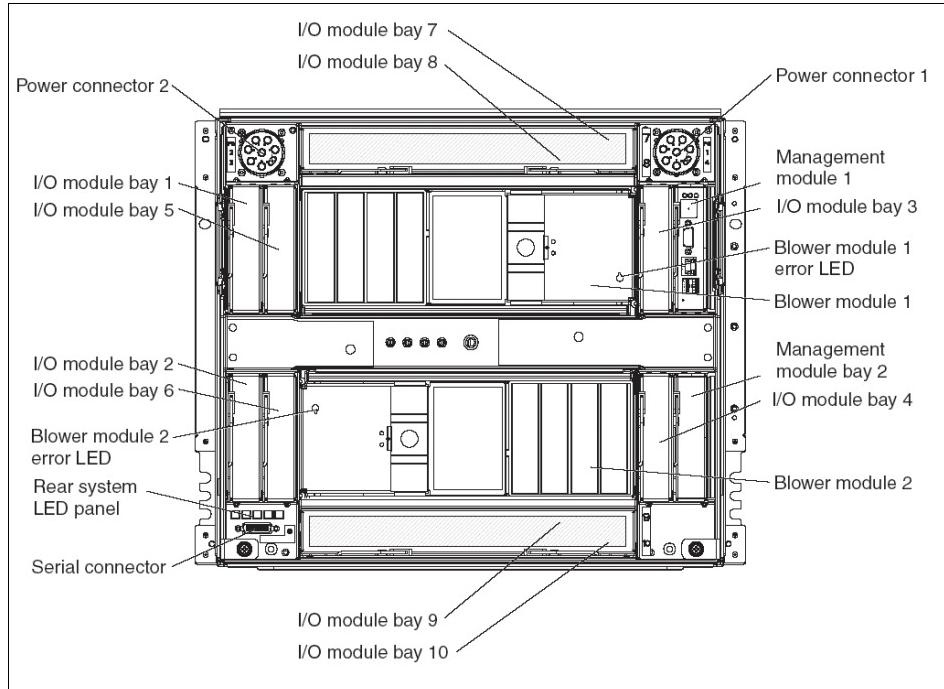


Figure 2-4 Rear view of BladeCenter H

The key features on the rear of the BladeCenter H are:

- ▶ Two hot-swap blower modules as standard
- ▶ Two hot-swap management module bays—with one management module as standard
- ▶ Four traditional fabric switch modules
- ▶ Four high-speed fabric switch modules

The BladeCenter H chassis allows for either 14 single-slot blade servers or seven double-slot blade servers. However, you can mix different blade server models in one chassis to meet your requirements.

The BladeCenter H chassis ships standard with one Advanced Management Module. This module provides the ability to manage the chassis as well as providing the local KVM function. The optional redundant Advanced Management Module provides the IBM BladeCenter H with higher levels of resiliency. While in the chassis, the second module is in passive or standby mode. If the active or primary module fails, the second module is automatically enabled with all of the configuration settings of the primary module. This function

provides clients with easy remote management and connectivity to the BladeCenter H chassis for their critical applications.

BladeCenter H does not ship standard with any I/O modules. You choose these I/O modules based on your connectivity needs. An Ethernet Switch Module (ESM) or Passthrough Module will be required in I/O module bays 1 and 2, to enable the use of both Ethernet ports on a blade server. The I/O modules required in I/O module bays 3 and 4 depend on the I/O Expansion Card installed in the blade servers. The I/O modules required in the high-speed I/O module bays 7, 8, 9, and 10 depend on the HSFF or CFFv I/O Expansion Adapter cards installed in the blade servers.

The BladeCenter H chassis comes standard with no power cord. Power cords need to be ordered separately.

The BladeCenter H chassis ships standard with:

- ▶ One advanced management module
- ▶ Two blower modules
- ▶ Two power supply modules (one pair of 2,900-watt power modules).
- ▶ Two hot-swap power supply fan modules
- ▶ Two USB v2.0 ports
- ▶ One DVD-ROM drive

The chassis does not have a diskette drive. An optional USB-attached 1.44 MB diskette drive is available. The feature code for ordering the external diskette drive is FC 2591 of FRU P/N 03N4962.

Table 2-22 on page 28 provides information on the BladeCenter H (BCH) specifications.

Table 2-22 BladeCenter H specifications

Feature	Specification
Machine type	8852-4XU
Rack dimension	9U x 28 inches (711 mm)
DVD/CD standard drive	1 x DVD ROM (in Media Tray)
Diskette drive	None
Number of blade slots	14 (30 mm blade servers)
Number of switch module slots	10 hot-swap (4 x high speed, 4 x standard, 2 x bridge)

Feature	Specification
Switch module standard	None (in standard chassis offerings)
Power supply	2900 W AC
Number of power supplies (standard/maximum)	2 / 4 ^a
Number of blowers (standard/maximum)	2 / 2
Dimensions	Height: 15.75 inch (400 mm) Width: 17.40 inch (422 mm) Depth: 28.00 inch (711 mm)

a. Four power supplies are required to use high-speed bays 7 to 10, and any blade server in slots 8 to 14.

2.6.2 BladeCenter S

The BladeCenter S chassis is a robust and flexible physical platform. Its modular tool-free design allows easy access and maintenance. All external components (except running blade servers) are hot swappable and release levers/handles are clearly marked.

Figure 2-5 on page 30 and Figure 2-6 on page 31 display the front and rear view of an IBM BladeCenter S.

The key features on the front of the BladeCenter S are:

- ▶ A media tray at the front, with a DVD drive, two USB 2.0 ports, a system status LED panel, and two bay locations that are reserved for future use.
- ▶ Six hot-swap blade server bays supporting different blade server types.
- ▶ Two bays for disk storage modules; each storage module can house up to six 3.5-inch disk drives of internal storage. No storage modules are standard with the BladeCenter S chassis.

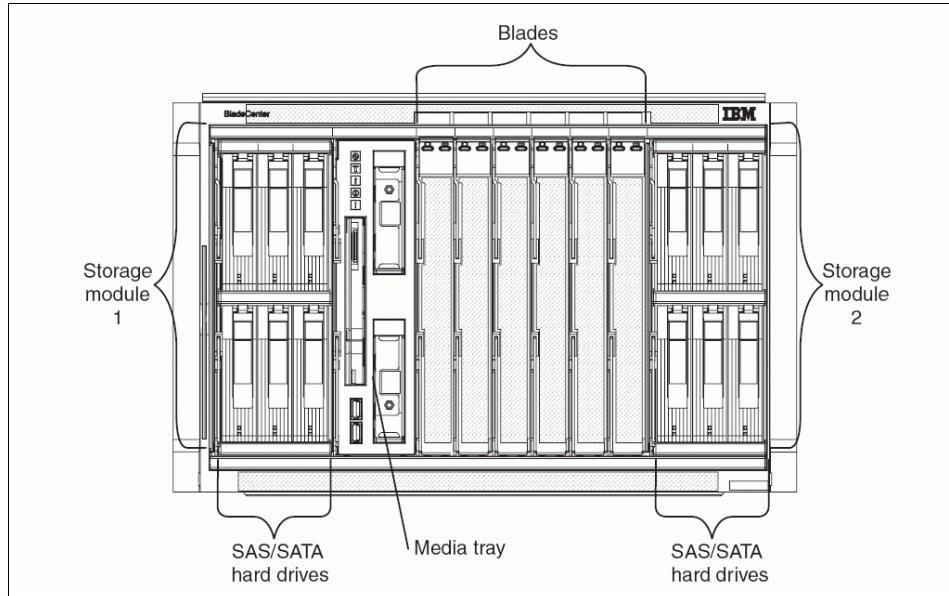


Figure 2-5 *BladeCenter S front view*

The key features on the rear of the BladeCenter S are:

- ▶ Four hot-swap blower modules as standard.
- ▶ One hot-swap management-module bay with one management module as standard.
- ▶ Four I/O bays for standard switch modules (bays 1, 3, and 4 can be used for installing I/O modules, bay 2 is reserved for future use).
- ▶ One pair of 950/1450-watt power modules. An additional power module option (configured in pairs of two 950/1450 W feature 4548 power modules) is available.
- ▶ One bay for a serial pass-through module to give a direct serial connection to installed blades (for those blades with the functionality).

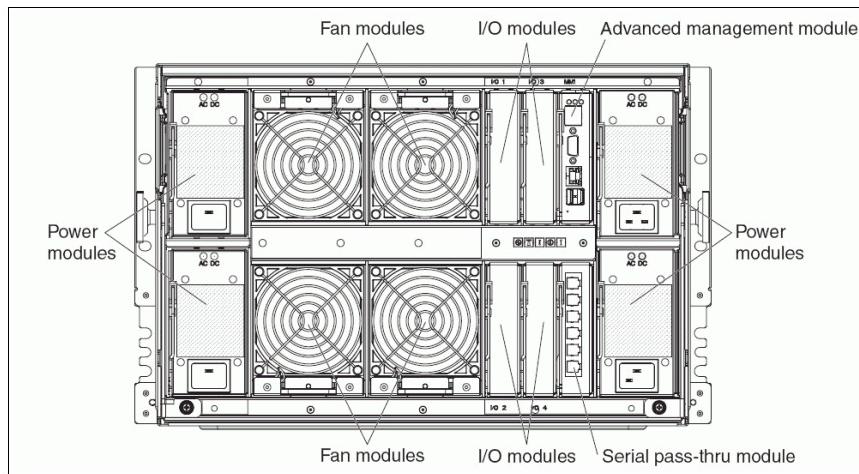


Figure 2-6 BladeCenter S rear view

The BladeCenter S chassis allows for either six single-slot blade servers or three double-slot blade servers. However, you can mix different blade server models in one chassis to meet your requirements.

The BladeCenter S chassis ships standard with an Advanced Management Module. This module provides the ability to manage the chassis as well as providing the local KVM function. Only one management module is supported with a BladeCenter S chassis.

The BladeCenter S chassis does not ship standard with any I/O modules. You choose these I/O modules based on your connectivity needs. An Ethernet Switch Module (ESM) is required in I/O module bay 1, to enable the use of both Ethernet ports on a blade server. The I/O modules required in I/O module bays 3 and 4 will

depend on the I/O Expansion Card installed in the blade servers. Bay 2 is reserved for future use. The chassis does not ship with any storage modules.

The BladeCenter S chassis uses either 100 to 127 v or 200 to 240 v AC power and can be attached to standard office power outlets.

The BladeCenter S chassis ships standard with:

- ▶ One advanced management module
- ▶ Four blower modules
- ▶ Two power supply modules (one pair of 950/1450-watt power modules)
- ▶ Two 2.5 m rack jumper power cords (IEC 320 C19 - C20)
- ▶ Four country-specific line cords
- ▶ One CD-RW/DVD-ROM drive

The chassis does not have a diskette drive. An optional USB-attached 1.44 MB diskette drive is available.

Table 2-23 on page 32 provides information on the BladeCenter S (BCS) specifications.

Table 2-23 BladeCenter S specifications

Feature	Specification
Machine type	8886-1MY
Rack dimension	7U x 28.9 inches (733.4 mm)
DVD/CD standard drive	1 x DVD ROM (in Media Tray)
Diskette drive	None
Number of blade slots	6 (30mm blade servers)
Number of switch module slots	4 hot-swap (3 x standard, 1 x standard for future use)
Switch modules (std/max)	0 / 4
Storage modules (std/max)	0 / 2
Power supply	950 W connected to 100 V AC 1450 W connected to 220 V AC
Number of power supplies (standard/maximum)	2 / 4

Feature	Specification
Number of blowers (standard/maximum)	4 / 4
Dimensions	Height: 12.00 inch (306.3 mm) Width: 17.50 inch (440 mm) Depth: 28.90 inch (733.4 mm)

2.6.3 BladeCenter HT

The IBM BladeCenter HT is a 12-server blade chassis designed for high-density server installations, typically for telecommunications use. It offers high performance with the support of 10 Gb Ethernet installations. This 12U high chassis with DC or AC power supplies provides a cost-effective, high-performance, high-availability solution for telecommunication network and other rugged non-telco environments. The IBM BladeCenter HT chassis is positioned for expansion, capacity, redundancy, and carrier-grade NEBS level/ETSI compliance in DC models.

IBM BladeCenter HT provides a solid foundation for next-generation networks (NGN), enabling service providers to become on demand providers. Coupled with technological expertise within the enterprise data center, IBM leverages the industry know-how of key Business Partners to jointly deliver added value within service provider networks.

The BladeCenter HT brings significant new capabilities to the broad IBM ecosystem of hundreds of NGN applications already being deployed on BladeCenter. A key example is the introduction of the Nortel 10 Gb Ethernet Switch Module for BladeCenter, which delivers 10 Gb to each blade server deployed in the BladeCenter H or BladeCenter HT chassis, and six 10 Gb Ethernet uplinks. This capability helps to greatly reduce the cost of implementing IPTV and other high bandwidth NGN applications.

The key features of the BladeCenter HT include:

- ▶ Support for up to 12 blade servers, compatible with the other chassis in the BladeCenter family
- ▶ Four standard I/O and four high-speed I/O module bays, compatible with the other chassis in the BladeCenter family
- ▶ A media tray at the front with light path diagnostics, two USB 2.0 ports, and optional compact flash memory module support
- ▶ Two hot-swap management module bays (one management module standard)

- ▶ Four hot-swap power-module bays (two power modules standard)
- ▶ New serial port for direct serial connection to installed blades
- ▶ Compliance with the NEBS 3 and ETSI core network specifications

Figure 2-7 and Figure 2-8 on page 35 show the front and rear view of the IBM BladeCenter HT.

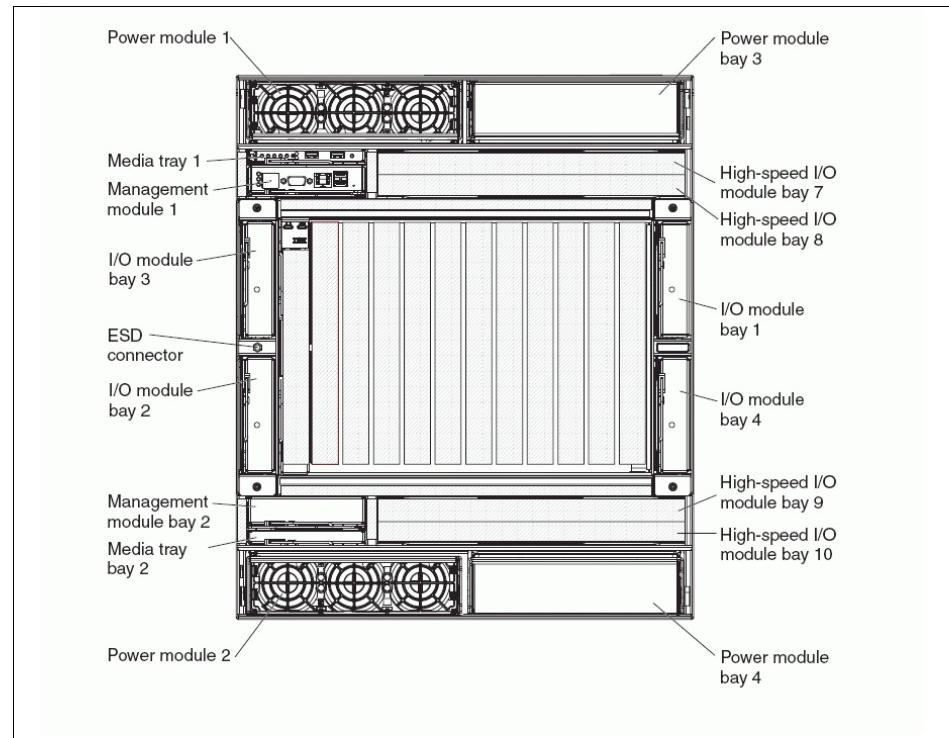


Figure 2-7 IBM BladeCenter HT front view

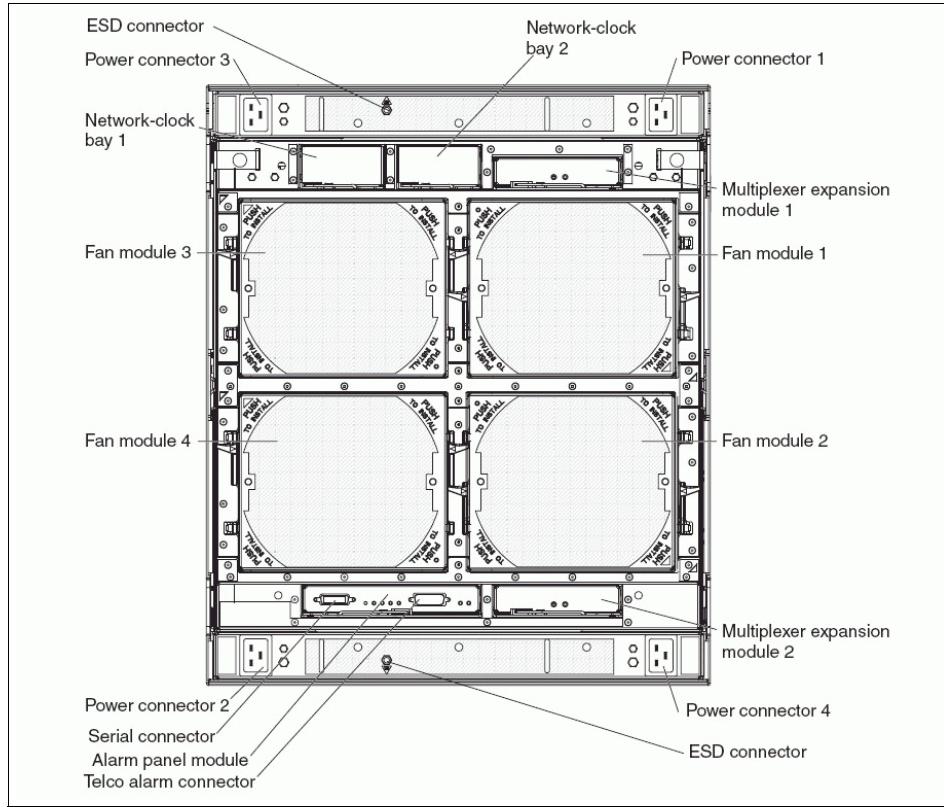


Figure 2-8 IBM BladeCenter HT rear view

Table 2-24 lists the features of the IBM BladeCenter HT.

Table 2-24 BladeCenter HT specifications

Feature	Specification
Machine type	8740-1RY (DC) 8750-1RY (AC)
Rack dimension	12U x 27.8 inches (706 mm)
DVD/CD standard drive	None
Diskette drive	None
Number of blade slots	12 (30mm blade servers)
Number of switch module slots	4

Feature	Specification
Number of high-speed switch module slots	4
Switch modules (std/max)	None
Number of power supplies (standard/maximum)	2 / 4 ^a
Number of blowers (standard/maximum)	4 / 4
Dimensions	Height: 21.00 inch (528 mm) Width: 17.50 inch (440 mm) Depth: 27.8 inch (706 mm)

- a. Four power supplies are required to use the high-speed bays 7 to 10, and any blade servers in slots 7 to 12.

The BladeCenter HT chassis allows for either 12 single-slot blade servers or six double-slot blade servers. However, you can mix different blade server models in one chassis. The BladeCenter HT chassis ships standard with one Advanced Management Module. This module provides the ability to manage the chassis, as well as providing the local KVM function. The optional redundant Advanced Management Module provides the IBM BladeCenter HT with higher levels of resiliency. This module provides clients with easy remote management and connectivity to the chassis.

The BladeCenter HT does not ship standard with any I/O modules. You need to choose these I/O modules depending on your connectivity needs. An Ethernet Switch Module (ESM) is required in I/O module bays 1 and 2 to enable the use of both Ethernet ports on a blade server. The I/O modules required in I/O module bays 3 and 4 depend on the I/O Expansion Card installed in the blade servers.

High-speed switch modules can be installed into I/O bays 7 through 10, and are used together with high-speed expansion cards installed into blade servers. I/O bays 1 through 4 can also be used for bridge modules.

For more information related to IBM BladeCenter technology and BladeCenter products, see *IBM BladeCenter Products and Technology*, SG24-7523.

2.6.4 Number of IBM BladeCenter JS23 and JS43 Express in Supported Blade Center Chassis

IBM BladeCenter JS23 and JS43 Express have their own power consumption characteristics. The amount of power requirements for each type of blade dictates the number of blades supported in each Blade Center chassis.

Table 2-25 shows the power consumption values for the IBM BladeCenter JS23 and JS43 Express, running with maximum configuration (8 x 8GB, HSDC, 1Xe and HDD on the JS23, and this same configuration on each JS43's planar).

Table 2-25 Power Consumption for IBM BladeCenter JS23 and JS43 Express

		Power Consumption (W)	
		Full Power	-30% P6 and L3, Static Power Save
JS23		400	334
JS43		777	654

Tables in Figure 2-9 and Figure 2-10 on page 38 show the maximum number of blades, running in the maximum configuration, per supported Blade Center chassis.

	BCS				BCH				BCH-T			
	Total of 6 slots				Total of 14 slots, 7 in each Power Domain (PD)				Total of 12 slots, 6 in each Power Domain (PD)			
	110VAC PS		220VAC PS		Only PD1 (No PD2)		PD1 and PD2		AC Power Supply		DC Power Supply	
	2PS	4PS	2PS	4PS	PD1	None	PD1	PD2	PD1	PD2	PD1	PD2
Fully Redundant without Performance Reduction	1	5	2	5	6	-	6	6	5	6	4	4
Redundant with Performance Reduction	1	6	2	6	7	-	7	7	6	6	5 (*)	5 (*)
Basic Power Mode (Max Power Capacity)	3	6	5	6	7	-	7	7	6	6	6	6
	(*) It is possible to select a special JS23 configuration such that all 12 blades can be supported in the Redundant with Performance Reduction mode on a BCH-T chassis. For instance, JS23 with one HSDC card, 4 x 8GB of memory and 1 HDD.											

Figure 2-9 Maximum number of IBM BladeCenter JS23 Express blades per IBM Blade Center chassis.

	BCS				BCH				BCH-T			
	Total of 6 slots				Total of 14 slots, 7 in each Power				Total of 12 slots, 6 in each Power			
	110VAC PS	220VAC PS	Only PD1 (No PD2)	PD1 and PD2	AC Power Supply	DC Power Supply	PD1	PD2	PD1	PD2	PD1	PD2
	2PS	4PS	2PS	4PS	PD1	None	PD1	PD2	PD1	PD2	PD1	PD2
Fully Redundant without Performance Reduction	1 JS23	2 JS43 + 1 JS23	2 JS43	2 JS43 + 1 JS23	3 JS43	-	3 JS43	3 JS43	2 JS43 + 1 JS23	3 JS43	2 JS43	2 JS43
Redundant with Performance Reduction	1 JS23	3 JS43	2 JS43	3 JS43	3 JS43 + 1 JS23	-	7 JS43		3 JS43	3 JS43	2 JS43 + 1 JS23 (*)	2 JS43 + 1 JS23 (*)
Basic Power Mode (Max Power Capacity)	1 JS43 + 1 JS23	3 JS43	2 JS43 + 1 JS23	3 JS43	3 JS43 + 1 JS23	-	7 JS43		3 JS43	3 JS43	3 JS43	3 JS43
	(*) It is possible to select a special JS43 configuration such that all 6 blades can be supported in the "Redundant with Performance Reduction" mode on a BCH-T chassis. For instance, JS43 with one HSDC card, 4 x 8GB of memory on each planar and 1 HDD.											

Figure 2-10 Maximum number of IBM BladeCenter JS23 and JS43 Express per IBM Blade Center chassis.

Note: It is recommended to use the Blade Center chassis with all power supplies installed to get full advantage of all performance enhancements made in the IBM BladeCenter JS23 and JS43 Express. This will also maximize the number of blades supported within the chassis.

2.6.5 IBM Director

Use IBM Director to perform network and system management tasks.

With IBM Director, a network administrator can perform the following tasks:

- View the hardware configuration of remote systems, in detail
- Monitor the usage and performance of critical components, such as microprocessors, disks, and memory
- Centrally manage individual or large groups of IBM and non-IBM x86-processor-based servers, desktop computers, workstations, and notebook computers on a variety of platforms

IBM Director provides a comprehensive entry-level workgroup hardware manager. It includes the following key features:

- Advanced self-management capabilities for maximum system availability.
- Multiple operating-system platform support, including Microsoft® Windows Server® 2003, Windows 2000 Server, Windows XP Professional, Windows Vista®, Red Hat Linux, SUSE Linux, VMware, Novell® NetWare, AIX, and IBM i operating system (formerly known as i5/OS).
- Support for IBM and non-IBM servers, desktop computers, workstations, and notebook computers.
- Support for systems-management industry standards.

- Integration into leading workgroup and enterprise systems-management environments.
- Ease of use, training, and setup.

IBM Director also provides an extensible platform that supports advanced server tools that are designed to reduce the total cost of managing and supporting networked systems. By deploying IBM Director, you can achieve reductions in ownership costs through the following benefits:

- Reduced downtime
- Increased productivity of IT personnel and users
- Reduced service and support costs
- Reduced power consumption



Technical description of the hardware architecture

IBM BladeCenter JS23 Express is a single wide blade, while the IBM BladeCenter JS43 Express is a double wide blade, consisting of the JS23's Base planar and a *Multiple Expansion Unit* planar (MPE). The MPE planar design is similar to the base planar, but with reduced functions. In this chapter we present the technical details of JS23's Base planar, highlighting the differences to the MPE planar as appropriate.

This chapter discusses the hardware architecture, including:

- ▶ “POWER6 processor” on page 43
- ▶ “Simultaneous Multi Threading (SMT)” on page 45
- ▶ “AltiVec (SIMD)” on page 46
- ▶ “IBM EnergyScale technology” on page 47
- ▶ “Flexible Support Processor” on page 49
- ▶ “Management Card” on page 50
- ▶ “Memory subsystem” on page 50
- ▶ “I/O subsystem” on page 53
- ▶ “PowerVM” on page 61
- ▶ “Operating system support” on page 62

- ▶ “Systems management” on page 64

3.1 POWER6 processor

The POWER6 processor capitalizes on the enhancements brought by the POWER5 processor.

Two of the enhancements of the POWER6 processor are the ability to do processor instruction retry and alternate processor recovery. This significantly reduces exposure to both hard (logic) and soft (transient) errors in the processor core.

- ▶ **Processor instruction retry**

Soft failures in the processor core are transient errors. When an error is encountered in the core, the POWER6 processor will first automatically retry the instruction. If the source of the error was truly transient, the instruction will succeed and the system will continue as before. On predecessor IBM systems, this error would have caused a checkstop.

- ▶ **Alternate processor retry**

Hard failures are more difficult, being true logical errors that are replicated each time the instruction is repeated. Retrying the instruction will not help in this situation because the instruction will continue to fail. Systems with POWER6 processors introduce the ability to extract the failing instruction from the faulty core and retry it elsewhere in the system, after which the failing core is dynamically deconfigured and called out for replacement. The entire process is transparent to the partition owning the failing instruction. Systems with POWER6 processors are designed to avoid what would have been a full system outage.

- ▶ **POWER6 single processor checkstopping**

Another major advancement in POWER6 processors is single processor checkstopping. A processor checkstop would result in a system checkstop. A new feature in the BladeCenter JS23 Express server is the ability to contain most processor checkstops to the partition that was using the processor at the time. This significantly reduces the probability of any one processor affecting total system availability.

- ▶ **POWER6 cache availability**

In the event that an uncorrectable error occurs in L2 cache, the system will be able to dynamically remove the offending line of cache without requiring a reboot. In addition, POWER6 utilizes an L1/L2 cache design and a write-through cache policy on all levels, helping to ensure that data is written to main memory as soon as possible. It can also make use of an additional 32MB chip-shared L3 cache.

- ▶ **POWER6 processor modules on IBM BladeCenter JS23 Express and JS43 Express**

IBM BladeCenter JS23 Express comes with 2 POWER6 processor modules (4-way), and IBM BladeCenter JS43 Express comes with two additional POWER6 modules (total 8-way).

Each POWER6 module is 4-way Dual Core Module (DCM), containing two 64bit 2-core POWER6 processors (4.2GHz) and one 32MB L3 cache.

Figure 3-1 shows a high-level view of the POWER6 module present in the JS23 and JS43 Express servers.

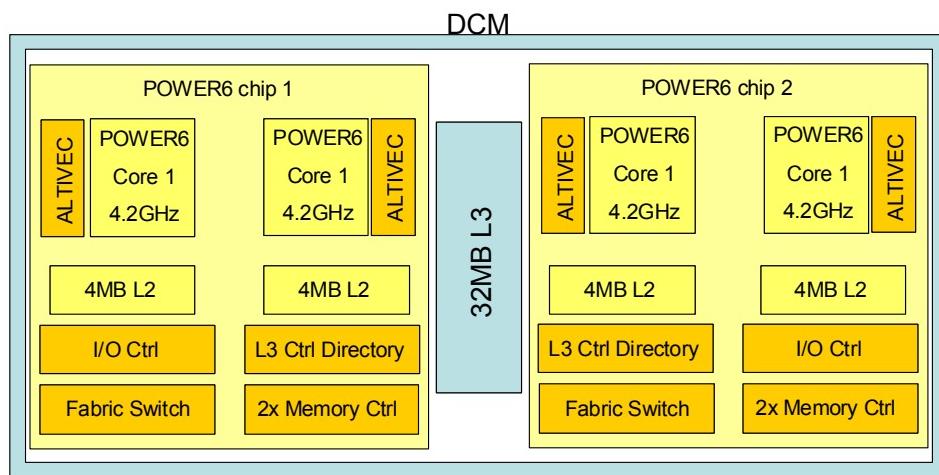


Figure 3-1 POWER6 processor module (4-way DMC)

3.1.1 Decimal floating point

This section describes the behavior of the decimal floating-point processor, the supported data types, formats, and classes, and the usage of registers. The decimal floating-point (DFP) processor shares the 32 floating-point registers (FPRs) and the floating-point status and control register (FPSCR) with the binary floating-point (BFP) processor. However, the interpretation of data formats in the FPRs, and the meaning of some control and status bits in the FPSCR are different between the BFP and DFP processors.

The DFP processor supports three DFP data formats: DFP32 (single precision), DFP64 (double precision), and DFP128 (quad precision). Most operations are performed on the DFP64 or DFP128 format directly. Support for DFP32 is limited to conversion to and from DFP64. For some operations, the DFP processor also

supports operands in other data types, including signed or unsigned binary fixed-point data, and signed or unsigned decimal data.

DFP instructions are provided to perform arithmetic, compare, test, quantum-adjustment, conversion, and format operations on operands held in FPRs or FPR pairs.

Arithmetic instructions These instructions perform addition, subtraction, multiplication, and division operations.

Compare instructions These instructions perform a comparison operation on the numerical value of two DFP operands.

Test instructions These instructions test the data class, the data group, the exponent, or the number of significant digits of a DFP operand.

Quantum-adjustment instructions

These instructions convert a DFP number to a result in the form that has the designated exponent, which may be explicitly or implicitly specified.

Conversion instructions These instructions perform conversion between different data formats or data types.

Format instructions These instructions facilitate composing or decomposing a DFP operand.

For example, the SAP® NetWeaver 7.10 ABAP™ kernel introduces a new SAP ABAP data type called DECFLOAT to enable more accurate and consistent results from decimal floating point computations. The decimal floating point (DFP) support by SAP NetWeaver® leverages the built-in DFP feature of POWER6 processors. This allows for highly simplified ABAP-coding while increasing numeric accuracy, and with a potential for significant performance improvements.

3.2 Simultaneous Multi Threading (SMT)

The POWER6 processor core has been designed to support both enhanced Simultaneous Multi Threading (SMT) and Single Threaded (ST) operation modes. Both IBM BladeCenter JS23 and JS43 Express supports the SMT technology.

Enhanced SMT features

To improve SMT performance for various workloads and provide robust quality of service, POWER6 provides two features:

- ▶ Dynamic resource balancing

The objective of dynamic resource balancing is to ensure that the two threads executing on the same processor flow smoothly through the system.

Depending on the situation, the POWER6 processor resource balancing logic has a different thread throttling mechanism (a thread reached threshold of L2 cache misses and will be throttled to allow other threads to pass the stalled thread).

- ▶ Adjustable thread priority

Adjustable thread priority lets software determine when one thread should have a greater (or lesser) share of execution resources. POWER6 supports eight software-controlled priority levels for each thread.

Single-threaded (ST) operation

Not all applications benefit from SMT. Having threads executing on the same processor does not increase the performance of applications with execution unit limited performance or applications that consume all the chip's memory bandwidth. For this reason, POWER6 supports the ST execution mode. In this mode, the POWER6 processor gives all the physical resources to the active thread. Highly optimized scientific codes are one example where ST operation is ideal.

3.3 AltiVec (SIMD)

IBM Semiconductor's advanced Single Instruction, Multiple Data (SIMD) technology based on the AltiVec instruction set is designed to enable exceptional general-purpose processing power for high-performance POWER processors. This leading-edge technology is engineered to support high-bandwidth data processing and algorithmic-intensive computations, all in a single-chip solution.

With its computing power, AltiVec technology also enables high-performance POWER processors to address markets and applications in which performance must be balanced with power consumption, system cost, and peripheral integration.

The AltiVec technology is a well-known environment for software developers who want to add efficiency and speed to their applications. A 128-bit vector execution unit was added to the architecture. This engine operates concurrently with the

existing integer and floating-point units and enables highly parallel operations, up to 16 operations in a single clock cycle. By leveraging AltiVec technology, developers can optimize applications to deliver acceleration in performance-driven, high-bandwidth computing.

The AltiVec technology is not comparable to the IBM POWER6 processor implementation, which uses the Simultaneous Multithreading functionality.

3.4 IBM EnergyScale technology

IBM EnergyScale technology is featured on the IBM POWER6 processor-based systems. It provides functions to help you understand and control IBM server power and cooling usage.

In this section we describe IBM EnergyScale features and hardware and software requirements.

Power Trending

EnergyScale provides continuous power usage data collection. This provides the administrators with the information to predict power consumption across their infrastructure and to react to business and processing needs. For example, an administrator could adjust server consumption to reduce electrical costs. To collect power data for the IBM BladeCenter JS23 Express you do not need any additional hardware because it collects the information internally.

Power Saver Mode

Power Saver Mode can be static or dynamic. Static Power Saver Mode lowers the processor frequency and voltage on a system a fixed amount, reducing the power consumption of the system while still delivering predictable performance. This percentage is predetermined to be within a safe operating limit and is not user configurable. Under current implementation this is a 14% frequency drop. When CPU utilization is low, Power Saver Mode has no impact on performance. Power Saver Mode can reduce the processor usage up to 30%. Power Saver Mode is not supported during boot or reboot although it is a persistent condition that will be sustained after the boot when the system starts executing instructions.

Dynamic Power Saver Mode, on the other hand, varies processor frequency and voltage based on the utilization of the system's POWER6 processors. When dynamic

power savings mode is enabled, the firmware of the system continuously monitors the utilization of the system, and adjusts the CPU clock speed and voltage to provide enough power to run the current workload. The less the system is utilized, the more power savings are achieved. In addition, you can specify whether you want to favor performance or favor power when enabling dynamic power savings mode. With favor performance, the peak frequency of the processors may be greater than 100%. With favor power, the processors are limited to 95% of nominal frequency under full utilization.

Power Capping

Capping enforces a user-specified limit on power usage. Power Capping is not a power saving mechanism. It enforces power caps by actually throttling the processors in the system, degrading performance significantly. The idea of a power cap is to set something that should never be reached but frees up margined power in the data center. The margined power is the amount of extra power that is allocated to a server during its installation in a datacenter. It is based on the server environmental specifications that usually are never reached because server specifications are always based on maximum configurations and worst case scenarios.

Processor Core Nap

The IBM POWER6 processor uses a low-power mode called Nap that stops processor execution when there is no work to do on that processor core, that is, both threads are idle. Nap mode allows the hardware to clock off most of the circuits inside the processor core. Reducing active power consumption by turning off the clocks allows the temperature to fall, which further reduces leakage (static) power of the circuits causing a cumulative effect. Unlicensed cores are kept in core Nap until they are licensed and return to core Nap whenever they are unlicensed again.

For more information regarding the EnergyScale Technology on POWER6 processor-based systems, see:
<http://www-03.ibm.com/systems/power/hardware/whitepapers/energyscale.html>

3.4.1 Thermal Power Management Device (TPMD)

The implementation of performance-aware power and thermal management for POWER6 processor-based systems is called the EnergyScale architecture, which meets a number of basic requirements for system-level power. IBM BladeCenter JS23 and JS43 Express implementation uses an integrated circuit called Thermal Power Management™ Device (TPMD), placed on the management card.

On IBM BladeCenter JS43 Express there is only one TPMD processor, located in the Base planar. This TPMD processor manages power consumption and thermal conditions of both Base and MPE planars.

The TPMD card can dynamically optimize the processor performance depending on processor power and system workload. The EnergyScale design supports a number of power and thermal management policies.

Benchmark Benchmark maximizes the single-threaded performance of the system by putting one core of each processor into a low-power state, as in POWER6 the nap mode.

Maximum performance

The EnergyScale implementation regulates the system in such a way as to use the maximum performance possible without violating the power or thermal limits of the system.

Power cap The system is set to use a previously defined power usage limit. Maximum power savings attempts to save as much power as possible for a given workload.

Optimal power/performance

In this mode the EnergyScale implementation changes the system to the most optimal power/performance settings on the basis of workload characteristics and the power and thermal environment.

3.5 Flexible Support Processor

Also known as the *Service Processor*, the *Flexible Support Processor* (FSP) is integrated on the system board. It provides control busses for all major components on the system. The FSP-1 also has a 128MB flash memory used to store a compressed image of the system's firmware image.

Note: The IBM BladeCenter JS43 Express has two Service Processors, one in the Base planar, and one in the MPE planar. The Service Processor located in the MPE planar has only I/O functions, and does not provide redundancy nor backup support to the FSP in the Base planar.

3.6 Management Card

The Management Card provides a mean for making the Anchor system information chip pluggable. Management Card's plug is located on Base planar, just below the DIMMs (see Figure 3-4 on page 54). The Management Card is *always present* in its specially designated connector on the Base planar.

Both IBM BladeCenter JS23 and JS43 Express have only one Management Card, located in the Base planar.

Management card contains the TPMD processor, RN50 video controller and associated SDRAM, RISCWatch, Async and debug ethernet ports, Anchor chip and P6 characterization connector, and it is a *Customer Replaceable Unit* (CRU).

3.6.1 Anchor chip on Management Card

The Anchor Smartchip resides on Management Card. This Anchor chip stores system-specific information. The pluggable Management Card makes it possible to transfer system-specific information from a faulty system CRU to the replacement CRU.

3.7 Memory subsystem

For IBM BladeCenter JS23 and JS43 Express, the memory controller is internal to the POWER6 processor. It is interfaced to memory buffer chips located on the system board. The memory buffer chip in the POWER6 processor-based server is the next generation Synchronous Memory Interface ASIC that connects memory controllers to the *Industry Standard Very Low Profile (VLP) Registered DIMMs (RDIMM)*.

3.7.1 Memory description of IBM BladeCenter JS23 and JS43 Express

IBM BladeCenter JS23 Express has two memory channels per POWER6 processor module (4 channel total), and each memory channel connects to a memory buffer chip. This same configuration is present on the MPE planar of a IBM BladeCenter JS43 Express, for a total of 8 channels.

Each memory buffer chip connects to two Registered DDIMMs, giving a total of 8 DIMMs in the IBM BladeCenter JS23 Express, and 16 DIMMs in the BladeCenter JS43 Express.

Figure 3-2 show the memory wiring layout for the base planar on an IBM BladeCenter JS23 Express. This same wiring layout is used in the MPE planar of an IBM BladeCenter JS43 Express.

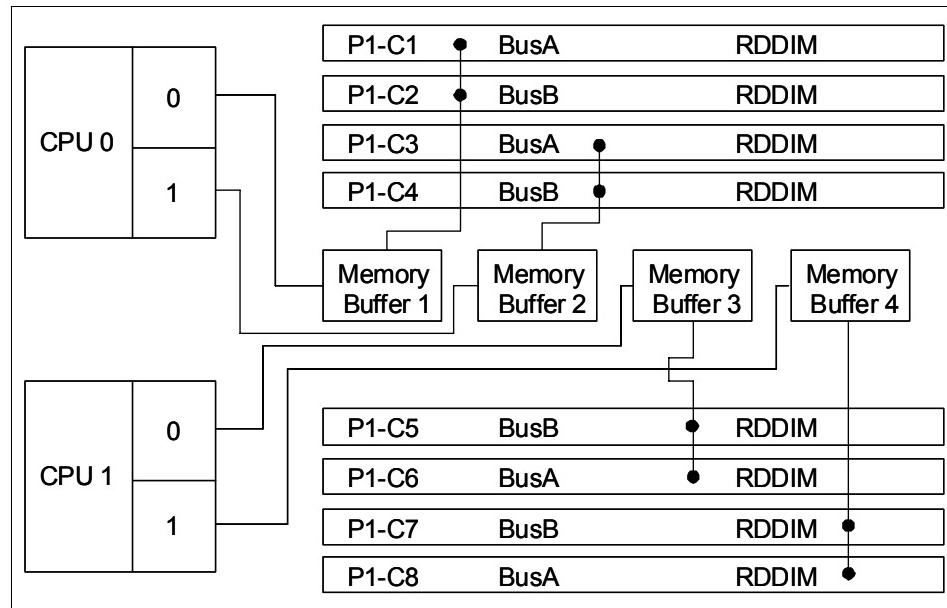


Figure 3-2 IBM BladeCenter JS23 Express memory wiring layout (Base planar - P1)

3.7.2 Memory placement rules for IBM BladeCenter JS23 and JS43

The following are the memory placement rules for IBM BladeCenter JS23 and JS43:

1. Install DIMM fillers in unused DIMM slots for proper cooling.

2. DDIMs are to be installed in pairs. First filling BusA then BusB of each planar, as shown above:
 - a. Base planar (P1): (C1, C3), (C6, C8), (C2, C4), (C5, C7).
 - b. MPE planar (P2): (C1, C3), (C6, C8), (C2, C4), (C5, C7).

Important: Both IBM BladeCenter JS23 and JS43 require a minimum of 4GB (2 x 2GB DIMM), and we recommend to plug them in slots P1-C1 and P1-C3 (BusA), as shown in Figure 3-3 on page 52

3. Both DDIMs in a pair must be of same size, speed, and technology. It is possible to mix compatible DDIMs from different manufacturers.
4. All DDIMs controlled by the same POWER6 processor module must be the same size and speed.

Important: Install only supported Registered DDIMs. For a complete list of supported Registered DIMMs, please see 2.1, “Overview of the JS23 and JS43 Express blade servers” on page 10. You can also find the most recent list of supported DDIMs in the IBM ServerProven® web site:

<http://www-03.ibm.com/servers/eserver/serverproven/compat/us/>

Figure 3-3 shows the recommended memory placement layout for both IBM BladeCenter JS23 and JS43 Express.

RDDIM Count	JS23 Base blade planar (P1) RDDIM slots								JS43 Expansion unit planar (P2) RDDIM slots							
	P1-C1	P1-C2	P1-C3	P1-C4	P1-C5	P1-C6	P1-C7	P1-C8	P2-C1	P2-C2	P2-C3	P2-C4	P2-C5	P2-C6	P2-C7	P2-C8
2	X		X													
4	X		X				X				X					
6	X	X	X	X			X			X						
8	X	X	X	X	X	X	X	X								
RDDIM Count	JS43 Base blade planar (P1) RDDIM slots								JS43 Expansion unit planar (P2) RDDIM slots							
	P1-C1	P1-C2	P1-C3	P1-C4	P1-C5	P1-C6	P1-C7	P1-C8	P2-C1	P2-C2	P2-C3	P2-C4	P2-C5	P2-C6	P2-C7	P2-C8
2	X		X													
4	X		X							X		X				
6	X		X			X		X	X		X					
8	X		X			X		X	X		X				X	X
10	X	X	X	X		X		X	X	X	X			X	X	
12	X	X	X	X		X		X	X	X	X	X		X	X	
14	X	X	X	X	X	X	X	X	X	X	X	X		X	X	
16	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X

Figure 3-3 Recommended memory placement for IBM BladeCenter JS23 and JS43 Express

3.7.3 Memory RAS

IBM BladeCenter JS23 and JS43 Express supports Memory Scrubbing, ECC, Chipkill Correction and Bit Steering. You can find more details about these and other POWER Systems RAS technologies in the following white papers:

- ▶ *IBM POWER Systems: Designed for Reliability.*
http://www-01.ibm.com/common/ssi/cgi-bin/ssialias?infotype=SA&subtype=WH&htmlfid=POW03019USEN&attachment=POW03019USEN.PDF&appname=STGE_PO_PO_USEN_WH
- ▶ *IBM POWER Systems: Designed for Availability.*
http://www-01.ibm.com/common/ssi/cgi-bin/ssialias?infotype=SA&subtype=WH&htmlfid=POW03020USEN&attachment=POW03020USEN.PDF&appname=STGE_PO_PO_USEN_WH
- ▶ *IBM POWER Systems: Designing and Implementing Serviceability.*
http://www-01.ibm.com/common/ssi/cgi-bin/ssialias?infotype=SA&subtype=WH&htmlfid=POW03021USEN&attachment=POW03021USEN.PDF&appname=STGE_PO_PO_USEN_WH

3.8 I/O subsystem

In this section we will discuss the I/O subsystem of JS23 and JS43 blades.

3.8.1 PCI express

Both IBM BladeCenter JS23 and JS43 Express has support for *PCI-Express* (PCI-e) cards only. Each planar has connectors for one *PCI-e High Speed Daughter Card* (HSDC), and one *PCI-e Combinational I/O vertical Expansion (CIOv) Card*.

PCI-e is the successor to Peripheral Component Interconnect (PCI) and PCI-x bus systems, realized by point-to-point implementation with the following general I/O interface standards:

- ▶ Point-to-point serial interconnect with packetized, layered protocol
- ▶ 2.6 Gbits per pin pair in each direction (speed scale in future generations)
- ▶ Dual simplex connection
- ▶ Scalable bus widths
- ▶ Embedded clocking technique using 8-bit/10-bit encoding
- ▶ Isochronous data transfer support

- ▶ Compatible with PCI at the software layers

3.8.2 I/O Expansion Cards

IBM BladeCenter JS43 Express have two HSDC 450-pin connectors, one in each planar, and two CIOv 160-pin connectors, one in each planar as well.

Figure 3-4 on page 54 shows how the HSDC and CIOv cards fit together inside the Base planar of an IBM BladeCenter JS23 Express.

Note: IBM BladeCenter JS23 and JS43 Express supports only *Combined Form Factor* (CFFe) High Speed Daughter Cards. Neither *Small Form Factor* (SFF) nor *Large Form Factor* (LFF) High Speed Daughter cards are supported.

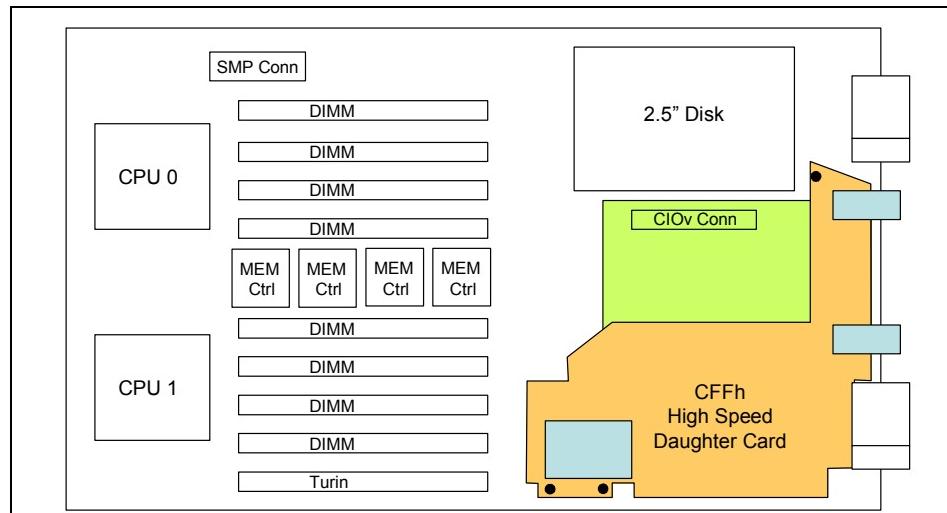


Figure 3-4 JS23 Base planar with a CFFh High Speed Daughter Card and CIOv Card and 2"5 HDD

Table 3-1 and Table 3-2 provides a list of some HSDC and CIOv daughter cards supported by the IBM BladeCenter JS23 and JS43 Express.

Table 3-1 Supported High Speed Daughter Cards (HSDC)

FRU Name	Feature	Supported OS
QLogic® 4Gb FChannel and Broadcom 1Gb Enet Combo	8252	AIX, Linux, IBM i

FRU Name	Feature	Supported OS
Mellanox 4X Infiniband Dual Port DDR Expansion Card	8258	AIX, Linux
Qlogic 8Gb FChannel	8271	Linux

Table 3-2 Supported CIOv PCI-e Expansion Cards

FRU Name	Feature	Supported OS
Emulex 8Gb Fibre Channel Expansion card	8240	AIX, Linux, IBM i
QLogic 4Gb FC Expansion Card	8241	AIX, Linux, IBM i
Qlogic 8Gb Fibre Channel Expansion card	8242	AIX, Linux, IBM i
3Gb SAS Passthrough Expansion Card	8246	AIX, Linux, IBM i

3.8.3 Universal Serial Bus (USB) subsystem

The USB bus connects USB devices to a USB host. The USB buses used in the BladeCenter chassis conform to the Universal Serial Bus 2.0 specification.

The IBM BladeCenter JS23 and JS43 Express uses the USB subsystem for shared resources such as keyboard, mouse, CD-ROM and floppy drive USB devices between the processor blades. These devices are connected to the processor blades via the Mid-plane.

The USB bus routed to the Management Module is used for the keyboard on the IBM BladeCenter JS23 and JS43 Express. This allows the Management Module to connect the keyboard to the IBM BladeCenter JS23 and JS43 Express, and control the routing of keyboard inputs to one of fourteen blades installed in a BladeCenter chassis.

The other USB bus is routed to the Media Tray, which connects to the DVD-ROM drive and diskette drive. Both USB busses are controlled by the Management Module independent of each other. This provides the capability to assign the CD and floppy disk drive to one blade while the keyboard is assigned to another blade. However, it does not restrict the assignment of shared resources of both USB buses to a single IBM BladeCenter JS23 or JS43 Express.

3.8.4 Integrated Virtual Ethernet (IVE)

IVE is the name given to the collection of hardware components (including the Host Ethernet Adapter (HEA), the software, and the hypervisor functions that provide the integrated Ethernet adapter ports with hardware assisted virtualization capabilities.

The IVE was developed to meet general market requirements for better performance and better virtualization for Ethernet. It offers:

- ▶ Two 1Gbps ports
- ▶ External network connectivity for LPARs using dedicated ports without the need of a Virtual I/O server
- ▶ The speed and performance of the GX+ bus, faster than PCI Express x16

The IVE consists of a physical Ethernet adapter that is connected directly to the GX+ bus of the POWER6 processor instead of being connected to a PCIe bus. This method provides IVE with the high throughput and low latency of a bus embedded in the I/O controller. IVE also includes special hardware features that provide logical Ethernet adapters.

The IVE feature implemented on the IBM BladeCenter JS23 Express provides two 1-Gbps physical ports. IBM BladeCenter JS43 Express has two additional 1-Gbps physical ports, located in the MPE planar.

Prior to IVE, virtual Ethernet provided a connection between LPARs. The use of an SEA and the Virtual I/O server allowed connection to an external network. The IVE replaces the need for both the virtual Ethernet and the SEA. It provides most of the functionality of each.

Therefore, this eliminates the need to move packets (using virtual Ethernet) between partitions and then through a shared Ethernet adapter (SEA) to an Ethernet port. LPARs can share IVE ports with improved performance.

The terms IVE and HEA are used interchangeably in this document.

3.8.5 Integrated Video Card

Both IBM BladeCenter JS23 and JS43 Express have an integrated ATI RN50 video controller in the Base planar. This video controller is used to support 2D graphics on an SVGA video display. This is the same controller as in the IBM BladeCenter JS21 Express, but with DDR2 SDRAM.

3.8.6 Serial Attached SCSI (SAS) storage subsystem

IBM BladeCenter JS23 and JS43 Express uses an embedded SAS controller that operates at 32-bit PCI-X at 133MHz.

Note: The SAS Drive in the JS23 Base planar is not hotpluggable.

On IBM BladeCenter JS23 Express there are four SAS ports. Two of them are wired to the SAS hard drive, and the other two go to the CIOv PCI-e connector, connecting to the Blade Center SAS Switch bay 3 and bay 4, when a SAS paddle card is used in the CIOv connector.

In Blade Center S, these two SAS ports are connected to the two *Data Storage Modules* (DSM) via the SAS Switch in bay 3 and bay 4. In Blade Center H, external storage are connected via the four external ports of the SAS Switches in bay 3 and bay 4.

IBM BladeCenter JS43 Express has four SAS ports that goes from the Base planar to the MPE planar, and they are used in similar functions as those on Base planar. IBM BladeCenter JS43 Express has only one SAS controller, located in the Base planar.

The boot HDD can either be on the Base or MPE planar of an IBM BladeCenter JS43 Express. The boot disk can also be from a Blade Center HDD via CIOv SAS paddle card.

Tip: It is recommended to install the boot HDD on the Base planar of the IBM BladeCenter JS43 Express.

SAS Drive

IBM BladeCenter JS23 and JS43 Express supports SAS 2.5" *Hard Disk Drive* (HDD) or SAS 2.5" *Solid Disk Drive* (SDD).

For a complete list of the supported SAS disks, please refer to Table 2-3 on page 12.

Important: SSD drivers are formatted in 528-Byte sector, and the SAS controller won't write to them in case it finds there are configuration sectors that are not clean. This situation happens when the SDD was previously used in RAID application. To be able to write to the SSD again, you will need to reformat it.

RAID support

IBM BladeCenter JS23 Express has no RAID available.

IBM BladeCenter JS43 Express has support for RAID functions when there are more than one SAS disk installed in the system. If there is only one drive then there is no RAID function.

For two drives in the IBM BladeCenter JS43 Express, the supported RAID functions are:

RAID 0 Striping.

RAID 1 Mirroring.

The drives on the Base Planar and MPE planar can be either rotating hard drives (HDD) or solid state drives (SSD). However, the SAS controller will not place an SSD and an HDD in the same RAID set.

Important: For RAID operations, all drives must be of the same type. Either HDD on both Base and MPE planars, or SDD on both Base and MPE planars.

Additionally, RAID can be done with different drives capacities, but you get only the effective capacity of the smaller device.

Configuring a SAS RAID array

Two SAS disk drives in the IBM BladeCenter JS43 Express can be used to implement and manage RAID level-0 and RAID level-1 arrays in operating systems that are on the ServerProven list.

For the blade server, you must configure the SAS RAID array through **smit sasdam**, which is the *SAS Disk Array Manager for AIX*.

The *SAS Disk Array Manager* is packaged with the Diagnostics utilities on the Diagnostics CD. Use the command **smit sasdam** to configure the disk drives for use with the SAS controller.

Tip: Refer to Chapter 3 of *IBM SAS RAID Controller for AIX* for more details on how to use the *IBM SAS Disk Array Manager*. This document is available at <http://publib.boulder.ibm.com/infocenter/systems/scope/hw/topic/arebj/arebj.pdf>

Important: Depending on your RAID configuration, you must create the array before you install the operating system in the blade server.

Before you can create a RAID array, you must reformat the hard disk drives so that the sector size of the drives changes from 512 MB to 528 MB. If you later decide to remove the hard disk drives, delete the RAID array before you remove the drives.

If you decide to delete the RAID array and reuse the hard disk drives, you must reformat the drives so that the sector size of the drives changes from 528 MB to 512 MB.

When the IBM BladeCenter JS43 Express is installed in a Blade Center Storage (BCS) and all or some of the BCS DASD are configured to that blade slots, those devices, if formatted to 528-Byte sector, can also be candidates for RAID array members with the JS43 Express drives themselves.

SAS CIOv paddle card

When the CIOv SAS paddle card is installed on an IBM BladeCenter JS23 or JS43 Express, it is possible to access the hard drives installed on a Blade Center S or external hard drive array.

Supported SAS topology via CIOv paddle card

When the CIOv SAS paddle card is installed on Base or MPE planar, the SAS controller can access the hard drives installed on a BCS or external hard drive array. The CIOv Paddle connects the SAS controller to the SAS switch bay 3 and bay 4. Each SAS Switch bay may be populated with a *Non-Raid SAS Switch Module* (NSSM).

Note: By the time of this writing, *RAID SAS Switch Modules* (RSSM) is not supported.

In BCS, the switch expands the SAS controller ports to the two internal DSMs and to storage enclosures that are attached to 4 external connections of the switch modules. In Blade Center H (BCH) or Blade Center HT (BCHT), only external storage enclosures are attached to the 4 external ports of the SAS switch. There is not internal storage in BCH nor BCHT.

Configurations are allowed with either one or both switch bays populated. Each NSSM contains 4 mini SAS connectors on it. The two SAS devices that can be attached to these connectors are SAS Tape devices and DS3200.

DS3200 is a RAID box that will appears as a SAS End device/Target with one or more LUNs. SAS Tape devices are also end devices.

Furthermore, when a IBM BladeCenter JS23 Express is installed in a BCS, there are “internal storage” Data Storage Modules in BCS that attaches to the NSSM

via the BladeCenter Mid planar. Each DSM has 2 SAS expanders with each expander connecting to the 6 DASD, one DSM connects to the primary ports of the DASD while the other expander connects to the secondary port of the DASD.

The A side expander of each DSM is wired to NSSM in switch bay 3 while the B side expander is wired to the NSSM in switch bay 4.

Figure 3-5 on page 60 shows the supported SAS topology for the IBM BladeCenter JS23 and JS43 Express on the BCS.

Figure 3-6 on page 61 shows the supported SAS topology for the IBM BladeCenter JS23 and JS43 Express on the BCH and BCHT.

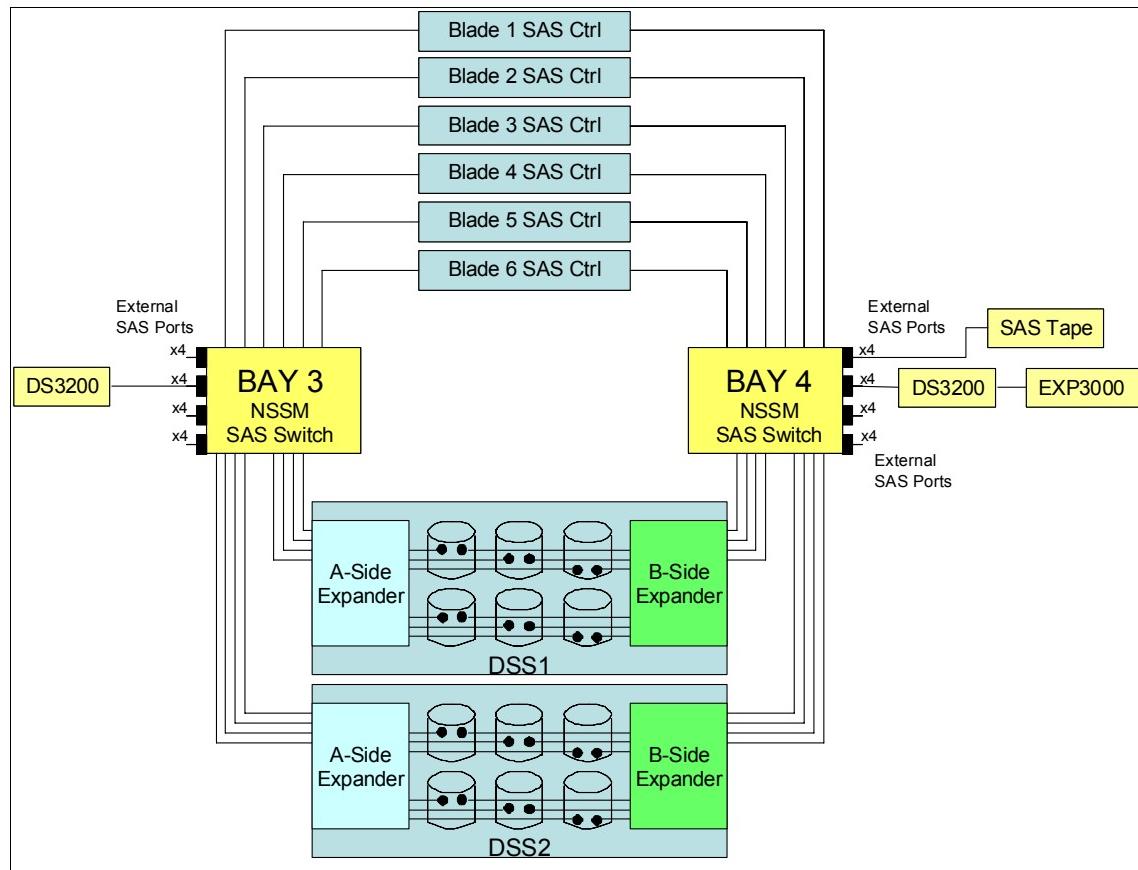


Figure 3-5 IBM BladeCenter JS23 and JS43 Express BCS SAS Topology

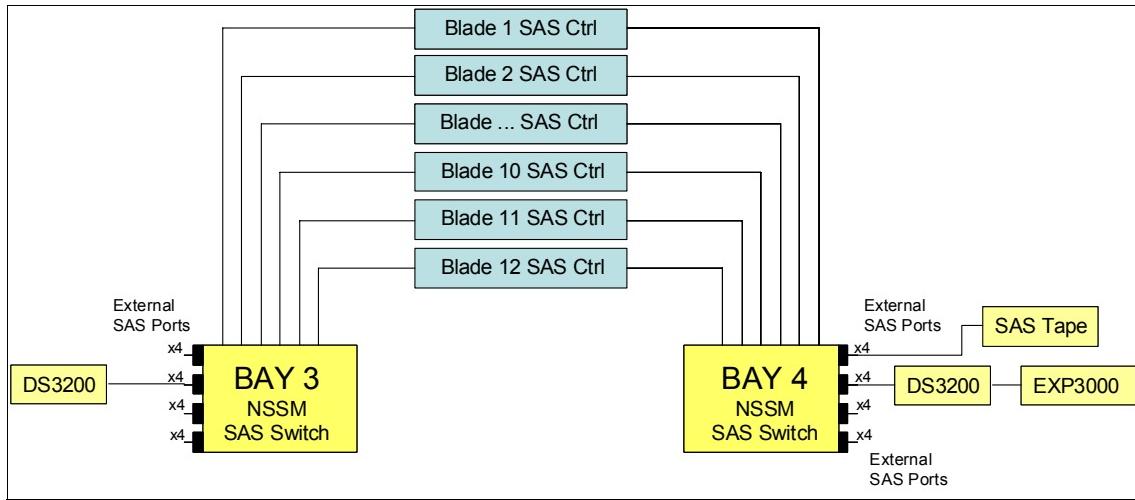


Figure 3-6 IBM BladeCenter JS23 and JS43 Express BCH and BCHT SAS Topology

3.9 PowerVM

The PowerVM platform is the family of technologies, capabilities, and offerings that provide virtualization capabilities on the IBM BladeCenter JS23 and JS43 Express systems. PowerVM is the new umbrella branding term for Power Systems Virtualization (Logical Partitioning, Micro-Partitioning™, Hypervisor, Virtual I/O Server, Advanced Power Virtualization, Live Partition Mobility, Workload Partitions, and so on).

As with Advanced Power Virtualization in the past, PowerVM is a combination of hardware enablement and value-added software. Table 3-3 on page 62 shows the licensed features of each of the two different editions of PowerVM for IBM BladeCenter JS23 and JS43 Express.

PowerVM standard edition offers a Web-based interface for managing virtualization within a single blade. The Integrated Virtualization Manager (IVM) component of Virtual I/O server (VIOS) provides the setup of management capabilities of logical partitions (LPARs). It manages Virtual I/O and Virtual Ethernet so that storage and communications adapters can be shared among all the LPARs running on the IBM BladeCenter JS23 and JS43 Express.

Upgrade from PowerVM Standard Edition to Enterprise Edition is possible and completely undisruptive. The upgrade does not even require the installation of additional software. Just enter a valid activation key code in the hypervisor to unlock the next level of functions.

Table 3-3 PowerVM editions for IBM BladeCenter JS23 and JS43 Express

Description	Standard Edition	Enterprise Edition
Maximum LPARs	10 / core	10 / core
Virtual I/O server	YES	YES
Integrated Virtualization Manager	YES	YES
Shared Dedicated Capacity	YES	YES
Live Partition Mobility	NO	YES
Active Memory Sharing	NO	YES

3.10 Operating system support

The IBM BladeCenter JS23 and JS43 Express supports the IBM AIX, Novell SUSE Linux Enterprise Server, Red Hat Enterprise Linux, and IBM i operating systems.

When using the PowerVM Standard Edition or Enterprise Edition on an IBM BladeCenter JS23 and JS43 Express running a VIOS, a combination of IBM i, IBM AIX, and Linux partitions can be deployed.

The following sections describe in detail the supported operating system versions and levels.

3.10.1 AIX

The following versions of AIX 5L™ are supported on the IBM BladeCenter JS23 and JS43 Express:

- ▶ AIX V5.3 with the 5300-07 Technology Level with Service Pack 9, or later
- ▶ AIX V5.3 with the 5300-08 Technology Level with Service Pack 7, or later
- ▶ AIX V5.3 with the 5300-09 Technology Level with Service Pack 4, or later
- ▶ AIX V5.3 with the 5300-10 Technology Level, or later
- ▶ AIX V6.1 with the 6100-03 Technology Level, or later
- ▶ AIX V6.1 with the 6100-02 Technology Level with Service Pack 4, or later
- ▶ AIX V6.1 with the 6100-01 Technology Level with Service Pack 5, or later
- ▶ AIX V6.1 with the 6100-00 Technology Level with Service Pack 9, or later

IBM periodically releases maintenance packages for the AIX 5L operating system. These packages are available on CD-ROM, or you can download them from:

<http://www.ibm.com/eserver/support/fixes/fixcentral/main/pseries/aix>

The Web page provides information about how to obtain the CD-ROM. You can also get individual operating system fixes and information about obtaining AIX 5L service at this site.

In AIX 5L V5.3, the **suma** command is also available, which helps the administrator to automate the task of checking and downloading operating system downloads. For more information about the suma command functionality, visit:

<http://www.ibm.com/systems/p/os/aix/whitepapers/suma.html>

3.10.2 Linux

The following versions of Linux for IBM POWER Series are supported on BladeCenter JS23 and JS43 Express:

- ▶ SUSE Linux Enterprise Server 10 Service Pack 2 for POWER Systems, or later
- ▶ SUSE Linux Enterprise Server 11 for POWER Systems, or later
- ▶ Red Hat Enterprise Linux 4.6 for POWER
- ▶ Red Hat Enterprise Linux 4.7 for POWER, or later
- ▶ Red Hat Enterprise Linux 5.1 for POWER
- ▶ Red Hat Enterprise Linux 5.2 for POWER
- ▶ Red Hat Enterprise Linux 5.3 for POWER, or later

For information about the features and external devices supported by Linux for IBM POWER on the IBM BladeCenter JS23 and JS43 Express server, visit:

<http://www-03.ibm.com/systems/bladecenter/hardware/servers/index.html>

For information about SUSE Linux Enterprise Server 10, visit:

<http://developer.novell.com/yesssearch/Search.jsp>

For information about Red Hat Enterprise Linux, visit:

<https://hardware.redhat.com/?pagename=hcl&view=certified&vendor=4&class=8>

Many of the features described in this document are operating system dependent and might not be available on Linux. For more information, visit:

http://www.ibm.com/systems/p/software/whitepapers/linux_overview.html

3.10.3 IBM i

IBM i 6.1 is supported on both IBM BladeCenter JS23 and JS43 Express. It uses IBM PowerVM Standard Edition, which includes the POWER Hypervisor™, Micro Partitioning, and Virtual I/O server with Integrated Virtualization Manager (IVM). PowerVM Standard Edition is available for no additional charge on the BladeCenter JS23 and JS43 Express.

For more information about IBM i operating systems running on IBM BladeCenter JS23 and IBM BladeCenter JS43 Express, see:

<http://www.ibm.com/systems/power/hardware/blades/ibmi.html>

3.11 Systems management

Several tools are available to manage IBM BladeCenter JS23 and JS43 Express. The BladeCenter Web interface can effectively manage the blades. In addition, IBM Director and Cluster Systems Management (CSM) are management tools that can be used to manage your environment.

3.11.1 BladeCenter Advanced Management Web interface

The BladeCenter Web interface allows system administrators to easily and effectively manage up to 14 blades from an integrated interface. From trivial tasks such as powering blades on or off, to more complex tasks such as firmware management, the Web interface allows powerful control over all blades and I/O modules that are attached to the BladeCenter chassis.

Important: IBM BladeCenter JS23 and JS43 Express does not support firmware updates through the Advanced Management Module. Their server enhanced service processor has a larger firmware image that makes it impractical to download and install over the RS-485 bus of the management module.

For more details on how to perform a firmware update, see 12.1, “Firmware updates” on page 460.

The BladeCenter Web interface allows the following:

- ▶ A System Administrator can easily and effectively manage up to 14 blade servers from an integrated interface.
- ▶ Power the IBM BladeCenter JS23 and JS43 Express on or off.
- ▶ Control over all blade servers and input/output (I/O) modules that are attached to the BladeCenter chassis even with a mixed environment.
- ▶ Manage other BladeCenter resources such as I/O modules and retrieval of system health information.
- ▶ Configure SoL for the IBM BladeCenter JS23 and JS43 Express.

3.11.2 IBM Director

IBM Director is a graphical user interface (GUI)-based management tool that provides management functions for the BladeCenter. IBM Director enables you to remotely manage many IBM and non-IBM servers, including the IBM BladeCenter JS23 and JS43 Express. The IBM Director console allows System Administrators to manage multiple BladeCenter chassis in a heterogeneous environment or environments where a Director infrastructure exists.

Visit the following Web site to download the IBM Director 6.1, and get the latest information on IBM Director 6.1 to IBM BladeCenter JS23 and JS43 Express:

<http://www.ibm.com/systems/management/director/downloads/>

IBM Systems Director Active Energy Manager

IBM Systems Director Active Energy Manager is the strategic power management interface for all IBM server products. From Active Energy Manager a user can display power trending information, display thermal information, enable/disable power savings modes or set a power cap. An administrator may connect a single instance of Active Energy Manager to multiple platforms, including IBM Blade Center chassis with one ore more IBM BladeCenter JS23 and JS43 Express blades.

As the strategic power management interface, Active Energy Manager supports the most comprehensive feature set of all the EnergyScale user interfaces, and is delivered as an extension to IBM Director.

Active Energy Manager has now been integrated into the Web-based interface of Director 6.1. The result is tighter integration of Active Energy Manager and IBM Director, eliminating the separately-installable Active Energy Manager console interface of previous Active Energy Manager releases.

In addition to a new web interface and the ability to install the IBM Director server on AIX, Active Energy Manager leverages Director 6.1 to provide the following new features:

- ▶ Dynamic Power Save for supported POWER6 servers
- ▶ Threshold-driven events for power, input temperature, and effective CPU speed
- ▶ Enhanced power policies for individual servers, and groups of servers
- ▶ Expanded power capping ranges via “soft” power caps on supported servers
- ▶ Complete Command Line Interface (CLI) support
- ▶ Support for additional hardware, including:
 - SynapSense Wireless Sensor Nodes
 - Non-IBM Power Distribution Units (PDUs) from Eaton and Raritan
 - Uninterruptible Power Supplies from Eaton
 - Monitoring of equipment managed by Emerson-Liebert’s SiteScan, including Computer Room Air Conditioning (CRAC) units, chillers, UPS’s, and PDUs. This includes getting real-time data on power consumption, as well as alerts when there are problems

For more information about the IBM Active Energy Manager, see the following resources:

- ▶ <http://www.ibm.com/systems/management/director/extensions/actengmrg.html>
- ▶ <http://www-03.ibm.com/systems/power/hardware/whitepapers/energyscale.html>

3.11.3 Cluster Systems Management (CSM)

IBM Cluster Systems Management (CSM) provides several useful functions to manage a cluster from a single point-of-control. These include resource monitoring, automated monitoring and operation, remote hardware control, remote command execution, security, configuration file management, parallel network installation, and diagnostics.

CSM V1.7 supports the following functions on the IBM BladeCenter JS23 and JS43 Express:

- ▶ Hardware control
- ▶ Install and update software on nodes
- ▶ Distributed command execution
- ▶ File synchronization across cluster

- ▶ Monitoring synchronization across cluster
- ▶ Monitoring and automated response
- ▶ Automatic security configuration
- ▶ Management of node groups (static and dynamic)
- ▶ Diagnostics tools

For more information about CSM, visit:

- ▶ <http://www-03.ibm.com/systems/clusters/software/csm/>
- ▶ <https://www14.software.ibm.com/webapp/set2/sas/f/lopdiags/home.html>



Part 2

System configuration and management

Having covered the basic system information and architecture in Part 1, we expand on that to include how to get BladeCenter JS23 and JS43 up and running the supported operating systems, and several other management-oriented topics.



System planning and configuration using VIOS with IVM

This chapter describes how to perform basic system planning prior to and configuration after you install Virtual Input/Output Server (VIOS). The configuration can be done by using the command line interface (CLI) and user interface (UI). The Web browser-based UI is an integral part of the Integrated Virtualization Manager (IVM) and is included in the VIOS.

Complete VIOS configuration and maintenance information can be found in *System i and System p Using the Virtual I/O Server*, which is available at this link:

http://publib.boulder.ibm.com/infocenter/iseries/v1r3s/en_US/info/iphb1/iphb1.pdf

The PowerVM Editions website also contains useful information:

<http://publib.boulder.ibm.com/infocenter/systems/scope/hw/index.jsp?topic=/arecu/arecukickoff.htm>

Additional IVM information can be found in *Integrated Virtualization Manager on IBM System p5*, REDP-4061

This chapter has the following sections:

- ▶ “Planning considerations” on page 73
- ▶ “VIOS system management using IVM” on page 83
- ▶ “First VIOS login” on page 86
- ▶ “First IVM connection” on page 93
- ▶ “VIOS network management and setup” on page 100
- ▶ “VIOS Storage Configuration and Management” on page 121
- ▶ “Partition configuration for Virtual I/O Client (VIOC)” on page 144
- ▶ “Console access and activating a partition” on page 166

4.1 Planning considerations

When planning your system environment for a IBM BladeCenter JS23 or JS43 a complete overview of the BladeCenter, blades, network and storage should be reviewed. Crafting an overall solution will help to eliminate expensive rework.

4.1.1 General considerations

We'll start with the general considerations.

BladeCenter

An understanding of how network and other I/O traffic is routed through the mid-plane of a BladeCenter chassis is needed to plan external network and storage connections. Figure 4-1 on page 74 and Figure 4-2 on page 75 show the connections between a blade's integrated ports and expansion card ports to the individual switch module bays.

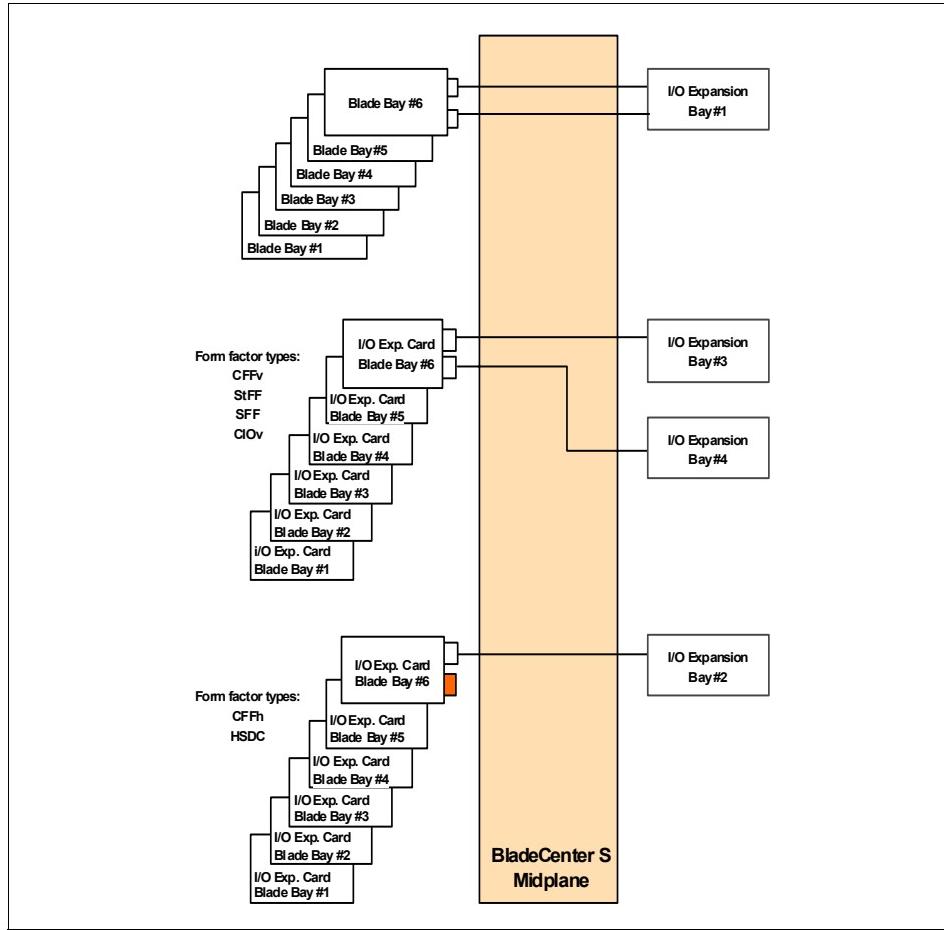


Figure 4-1 Expansion card port path through BladeCenter S midplane

Note: An RPQ is required when a Qlogic Ethernet and 4Gb Fibre Channel “combo card” CFFh expansion card is used in a BladeCenter S, only the two Ethernet ports are directed to the switch module in I/O Bay 2.

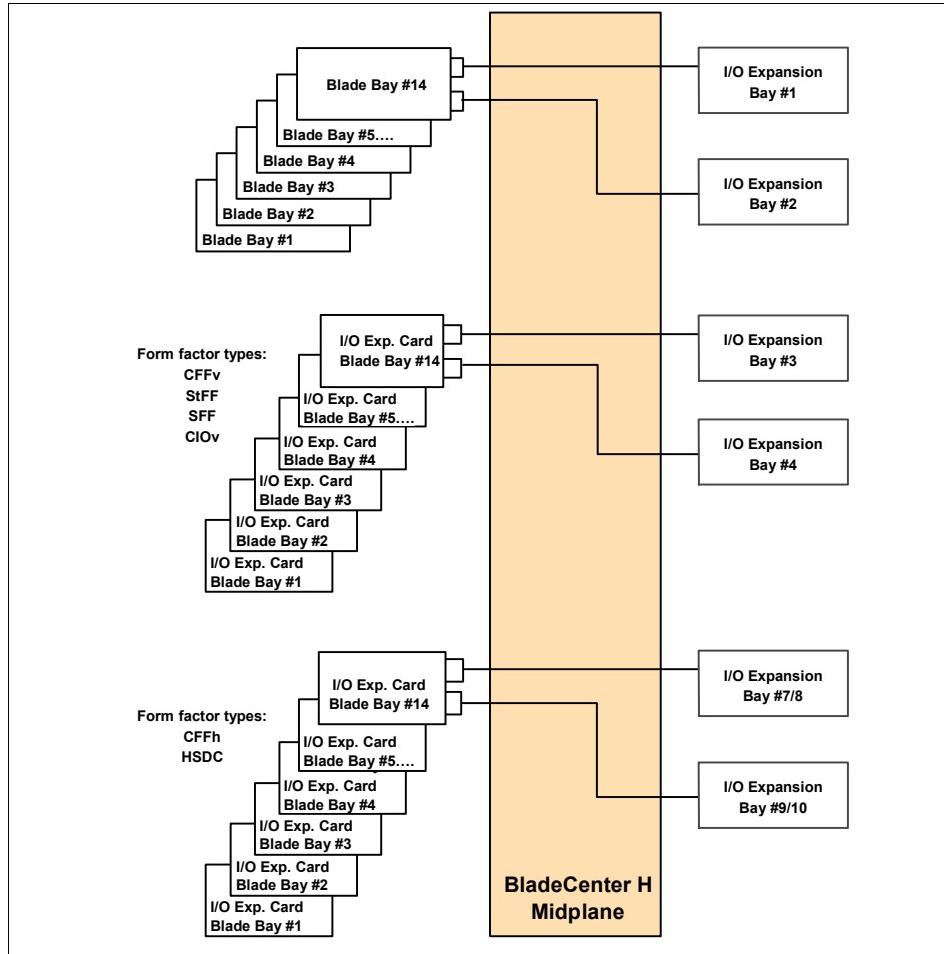


Figure 4-2 Expansion card port path through BladeCenter H mid plane

Internal storage

The IBM BladeCenter JS23 can have a single internal SAS drive. The IBM BladeCenter JS43 can have two internal drives. Disk mirroring is highly recommended for the operating system. The mirror can be between the internal drive and a SAN LUN, between two SAN LUNs, or between two internal drives when available.

Virtualization

After the physical connection planning aspects, the logical partitions or LPARs (assuming PowerVM virtualization) must be designed so that their requirements fit the CPU and memory available on the blade. Additional considerations are

required if you desire to use advanced operations available under PowerVM Enterprise Edition, such as Live Partition Mobility (LPM) and Active Memory™ Sharing (AMS).

The decision regarding whether to use a shared processor pool or dedicated processors should be made prior to configuring an LPAR. Changing from one mode to the other with the IVM UI requires the deletion of the LPAR and the creation of a new one, the VIOS CLI can use the **chsyscfg** command. The **chsyscfg** command can be used to switch when the partition is in the Not Activated state.

Network

EtherChannel configuration (if used), as well as BladeCenter switch module type and bay location, require pre-planning. HEA ports cannot be mixed with the PCI-X adapter ports available on the CFFh combo cards to create an EtherChannel device. The use of the HEA ports to create an EtherChannel will require the use of pass -thru modules, the Intelligent Copper Pass-thru Module for IBM BladeCenter in switch module bay one to retain Serial over LAN (SOL) capability, or Ethernet modules that provide for logical grouping of switches. Network Interface Backup or NIB, can mix HEA and other adapter port types

Note: The HEA adapter implementation on IBM BladeCenter JS23 or JS43 blades will always show a link status of Up, as shown using the **entstat** command:

```
$ entstat -all ent1 |grep Link  
Logical Port Link State: Up  
Physical Port Link State: Up
```

When an HEA adapter is used as primary in a network interface backup (NIB) configuration, the link state cannot be used as the failover determination. NIB should be set up using the *ping* option.

4.1.2 Internal and external storage considerations

Different types of storage can be used with the Virtual I/O Server. Each type of storage offers advantages and disadvantages. This section explains possible configuration scenarios with the IBM BladeCenter JS23 or JS43. For information about supported I/O expansion cards, along with their feature codes and option part numbers, refer to 3.8.2, “I/O Expansion Cards” on page 54.

There are currently two form factor types of I/O expansion cards supported in the IBM BladeCenter JS23 or JS43. They are:

- ▶ Combinational I/O vertical (CIOv) I/O expansion cards

- ▶ Combined Form Factor horizontal (CFFh) I/O expansion cards

CIOv adapter cards ports are always connecting to bay three and four of a BladeCenter chassis when installed in an IBM BladeCenter JS23 or JS43.

Figure 4-3 on page 77 show an Active SAS Pass through “paddle” expansion card in CIOv form factor. A QLogic 4 Gb Fibre Channel HBA, and Qlogic and Emulex 8Gb Fibre Channel HBAs are also available in the same form factor.

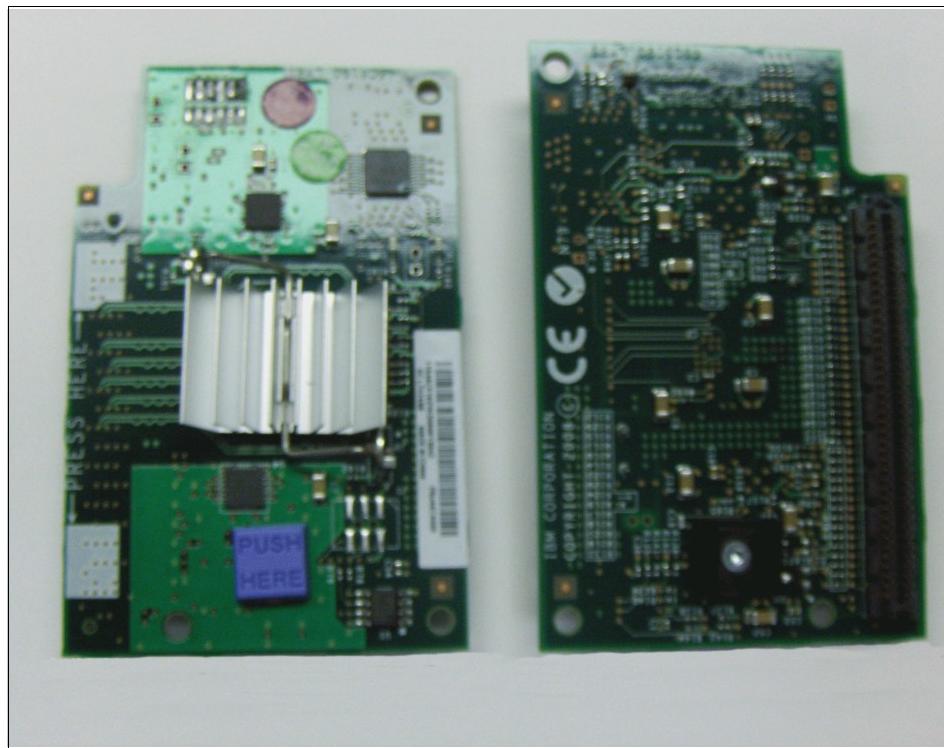


Figure 4-3 Active SAS Pass-through Expansion Card CIOv

The CFFh combo card has a dual port 4 Gb Fibre Channel Qlogic host bus adapter and a dual port Broadcom gigabit Ethernet adapter over a PCI-x bridge connected to the PCI-e bus of the blade. Figure 4-4 on page 78 shows the CFFh form factor. A Qlogic 8Gb Fibre Channel 2 port HBA in the CFFh format is also available.

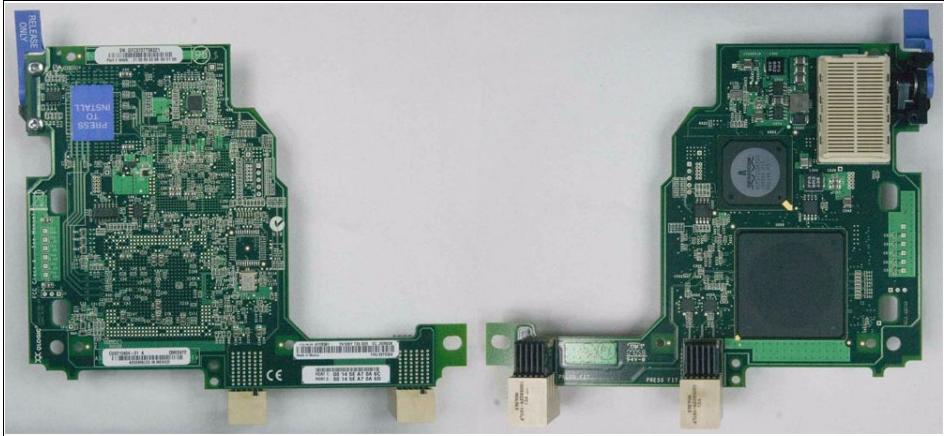


Figure 4-4 Qlogic Ethernet and 4 Gb Fibre Channel “combo card”CFFhTh

Together with an installed Qlogic Ethernet and 4Gb Fibre Channel combo card, it is also possible to install the CIOv I/O expansion card. Using a BladeCenter H with a JS23 combination gives, in addition to the two onboard network ports, six more I/O ports. These six additional ports are four Fibre Channel ports and two 1Gb Ethernet ports. The JS43 doubles this port count to four integrated network ports and twelve additional I/O ports or eight Fibre Channel ports and four additional 1Gb Ethernet ports.

Note: When a CIOv card and a CFFh card are installed together, the CIOv must be installed first.

You must install the matching type of I/O module in the bays for the type of I/O expansion card. Specifically, this means:

- ▶ You must install a module that provides Ethernet connectivity in module Bay 1 and Bay 2 of a BladeCenter H or Bay 1 of a BladeCenter S¹.
- ▶ You must install a module that supports the I/O connectivity of the installed CIOv expansion cards in Bay 3 and Bay 4.
- ▶ You must install a module or modules in a BladeCenter H Bays 7-10 or switch module in Bay 2 of a BladeCenter S that support the CFFh expansion cards.

Note: When the Optical Pass-thru Module is used with a 4 Gbs Fibre Channel host bus adapter, only a transfer rate of 2 Gbs can be used. Therefore, set the Fibre Channel switch ports that connect to the OPM to 2 Gbs.

¹ Both integrated network ports of a blade connect to the module in Bay 1. A VLAN configuration is required to separate the different networks and broadcast domains.

When JS23/43 blades with CFFh cards are installed in a BladeCenter H or HT the cards connect to the high speed bays 7, 8, 9, and 10 depending on the ports on the card. These module bays have a horizontal orientation. (The standard module bays have a vertical orientation.)

When JS23/43 blades with a supported CFFh card are installed in a BladeCenter S the cards are connected to Bay 2

Some CFFh cards utilize the high speed bays, but uses standard modules for connectivity. This connectivity is achieved by using the Multi-Switch Interconnect Module (MSIM). One MSIM² gets installed in Bays 7/8 and one in Bays 9/10. Each MSIM can be populated with a standard Fibre Channel switch module and an Ethernet switch module for the BladeCenter.

Figure 4-5 on page 79 shows a diagram of the MSIM and the type of switches that can be installed in the module bays. Refer to *BladeCenter Interoperability Guide* for details about which switches are supported in the MSIM. The guide can be found at:

<http://www-304.ibm.com/systems/support/supportsite.wss/docdisplay?lnocid=MIGR-5073016&brandind=5000020>

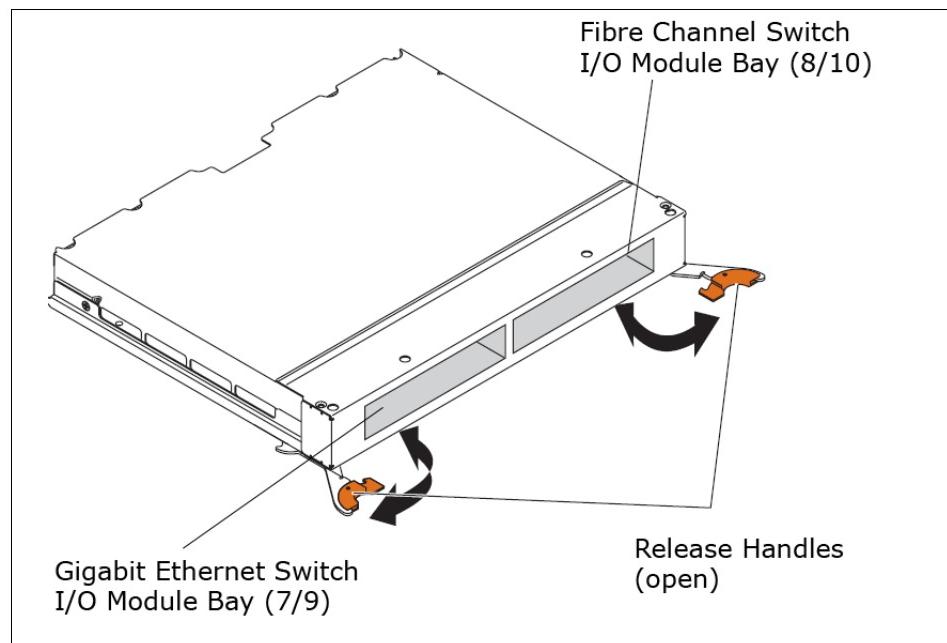


Figure 4-5 Multi-Switch Interconnect Module

² A interposer card is required for each MSIM to install in a BladeCenter HT chassis. Contact your IBM sales representative for information about ordering interposer cards.

JS23/JS43 storage

There are currently four different types of storage available:

- ▶ Internal 73GB or 146GB SAS Hard Disk Drive (HDD) disk storage
- ▶ Internal 73GB SAS Solid State Drive (SSD) storage
- ▶ External SAS/SATA disk storage
- ▶ External Fibre Channel storage

There is not a hardware initiator or TOE card available for the IBM BladeCenter JS23 or JS43 for iSCSI storage system attachment. Software initiators are available for AIX and Linux (no VIOS support).

Internal SAS storage

For internal storage, the JS23 blade can have one SAS HDD or SSD. The JS43 can have up to two SAS HDD or SSD storage devices installed.available. The JS43 supports RAID 0/1 types but both SAS devices must be of the same type HDD or SSD.

Note: Internal SAS drives are not currently supported for IBM i under the VIOS; in this case, they are only for VIOS use.

External SAS/SATA storage

External SAS or SATA storage is currently available in the BladeCenter S chassis. A BladeCenter S can be equipped with up to two Disk Storage Modules (DSM), which are each capable of holding up to six SAS or SATA hard drives.

This storage can be utilized by installing an Active SAS Pass through Expansion Card for IBM BladeCenter in the JS23 or JS43 and, one or two SAS Connectivity modules for non-RAIDed configurations. Two SAS RAID Controller Modules are required for RAIDed configurations.

With this setup, you are limited to the capacity that can be provided the two DSMs holding up to the maximum of 12 disks. More external storage could be made available using iSCSI and software initiators, but this is not recommended due to limited network port availability. You have only the two onboard network interfaces of the blade available.

External Fibre Channel storage

Various IBM and non-IBM Fibre Channel storage subsystems can be used to connect to a JS23 or JS43 installed in a BladeCenter, including models from the following series:

- ▶ IBM Total Storage DS8000™ series
- ▶ IBM Total Storage DS6000™ series
- ▶ IBM Total Storage DS5000™ series

- ▶ IBM Total Storage DS4000™ series
- ▶ IBM Total Storage DS3000™ series
- ▶ IBM Total Storage N™ series

The Virtual I/O Server data sheet provides an overview of supported storage subsystems and the failover driver that is supported with the subsystem. The data sheet can be found at:

http://www14.software.ibm.com/webapp/set2/sas/f/vios/documentation/data_sheet.html

Verify that your intended operating system supports these storage subsystems.

Also decide which type of I/O expansion card to use, and plan your storage area network. IBM offers Fibre Channel and Ethernet Switch modules from various vendors, and it is always easier to connect switches from the same vendor. In a Fibre Channel SAN environment, there will be zoning limitations if different switch vendors are used.

General storage considerations and support matrixes

You need to check a few support matrixes to plan your live partition mobility installation. This section points to support matrixes you can use to build a fully supported solution.

1. Start with the blade itself. All supported hardware and operating systems are listed on IBM ServerProven. ServerProven can be found at:

<http://www-03.ibm.com/servers/eserver/serverproven/compat/us/>

Click the middle of the page on the BladeCenter picture for blade- and BladeCenter-related information. You can find consolidated information from different sources in *BladeCenter Interoperability Guide* at:

<https://www-304.ibm.com/systems/support/supportsite.wss/docdisplay?lnocid=MIGR-5073016&brandind=5000020>

2. Verify the supported operating systems on the blade by using the link to NOS Support on the Compatibility for BladeCenter products page of ServerProven:
<http://www-03.ibm.com/servers/eserver/serverproven/compat/us/eserver.html>
3. Select the blade from the Compatibility for BladeCenter products page that you opened in step 1. A list of supported and tested hardware in combination with the blade itself is shown. Clicking a listed component retrieves details about the tested hardware and limitations that might exist for the component.
4. Using the information you gained about supported HBAs and storage products, you can start checking the Virtual I/O Server Data sheet to verify

which components supported by the blade are supported by the Virtual IO Server as well. The data sheet can be found at:

<http://www14.software.ibm.com/webapp/set2/sas/f/vios/documentation/datasheet.html>

5. Check the support matrix of the storage subsystem of your choice. In the case of Fibre Channel attached storage, verify the SAN switch support matrix. The following list points to the matrices of IBM storage products.

The DS8000 interoperability matrix can be found at:

<http://www.ibm.com/servers/storage/disk/ds8000/interop.html>

The DS6000 interoperability matrix can be found at:

<http://www.ibm.com/servers/storage/disk/ds6000/interop.html>

The DS5000 interoperability matrix can be found at:

<http://www-03.ibm.com/systems/storage/disk/ds5000/interop-matrix.html>

The DS4000 interoperability matrix can be found at:

<http://www.ibm.com/servers/storage/disk/ds4000/interop-matrix.html>

The DS3000 interoperability matrix can be found at:

<http://www-03.ibm.com/systems/storage/disk/ds3000/pdf/interop.pdf>

The ESS interoperability matrix can be found at:

<http://www.ibm.com/servers/storage/disk/ess/interop-matrix.html>

The N series interoperability matrix can be found at:

<http://www-03.ibm.com/systems/storage/nas/interophome.html>

The SAN volume Controller support matrix can be found at:

<http://www-304.ibm.com/jct01004c/systems/support/supportsite.wss/supportresources?taskkind=3&brandind=5000033&familyind=5329743>

The SAN switch interoperability matrix can be found at:

<http://www-03.ibm.com/systems/storage/san/index.html>

6. If you plan to implement IBM i on a JS23 or JS43, then in addition to the previous steps, verify the support matrix of IBM i on blades. Some restrictions apply when IBM i is used.

<http://www-03.ibm.com/systems/power/hardware/blades/ibmi.html>

http://www-03.ibm.com/systems/power/hardware/blades/supported_environments.pdf

7. After verifying the supported hardware components, check the storage area network cabling that is required for the storage product you plan to implement.

The cabling is described in the product documentation of the storage subsystem. Verify which failover drivers are supported by the storage subsystem. In the product documentation, check the recommended zoning configuration.

8. Use the Virtual I/O Server data sheet again to check which failover drivers are included in the Virtual I/O Server and which failover drivers can be installed.

Note: The System Storage™ Interoperation Center (SSIC) helps to identify supported storage environments. You find this Web-based tool at:

<http://www-03.ibm.com/systems/support/storage/config/ssic>

9. The storage subsystem usually requires a specific FCode, firmware, and driver level that is used with the selected host bus adapter in the blade. This information is not for every subsystem in the same location.

The Host bus adapters (HBA) Web-based tool helps to identify the required firmware level. You can check also product documentation or host system attachment guides when available. DS3000 and DS4000 support Web pages usually provide update packages for the supported adapters that contain the settings required for the HBA. They can be found at:

<http://www-03.ibm.com/systems/support/storage/config/hba/index.wss>

10. Define the boot device that you plan to use. You may boot from internal disk or external disk. Depending on the type of blade and I/O expansion card you plan to install, you may not have internal disks available. Redundancy may be not available at boot time for a boot disk. Manual interaction may be required and an administrator may need to be trained for situations where interaction with the system is required.

4.2 VIOS system management using IVM

Using VIOS is required when two or more partitions are required, and when using IBM i.

IVM provides a unique environment to administer logical partition-capable servers. It provides two ways to configure and manage logical partitions (LPARs):

- ▶ A user interface (UI) designed to be as simple and intuitive as possible, incorporating partition management, storage management, serviceability, and monitoring capabilities
- ▶ A command line interface (CLI), to enable scripting capabilities

You can use either interface to create, delete, and update the logical partitions and perform dynamic operations on LPARs (DLPAR) including the VIOS itself.

4.2.1 VIOS installation considerations

The Virtual I/O Server installation is performed like a native install of AIX. The basic requirements are:

- ▶ Console access, provided by Serial over LAN (SoL) or KVM
- ▶ AIX media or a NIM server
- ▶ A BladeCenter media tray assigned to the target blade, or network access to a NIM server
- ▶ Internal or external storage availability

4.2.2 IVM user interface

The IVM user interface is a Hypertext Markup Language (HTML)-based interface. It enables you to create LPARs on a single managed system, manage the virtual storage and virtual Ethernet on the managed system, perform maintenance, and view service information related to the managed system.

The IVM UI consists of several elements. The following two elements will be used most frequently:

Navigation area	The navigation area on the left side of the window displays the tasks and links that you can access in the work area.
Work area	The work area on the right side of the window contains information related to management tasks that you perform using IVM and related to the objects on which you can perform management tasks.

Figure 4-6 on page 85 shows the navigation and works of the IVM UI. UI help can be obtained from any window by clicking the ? link in the upper right corner of any window or view.

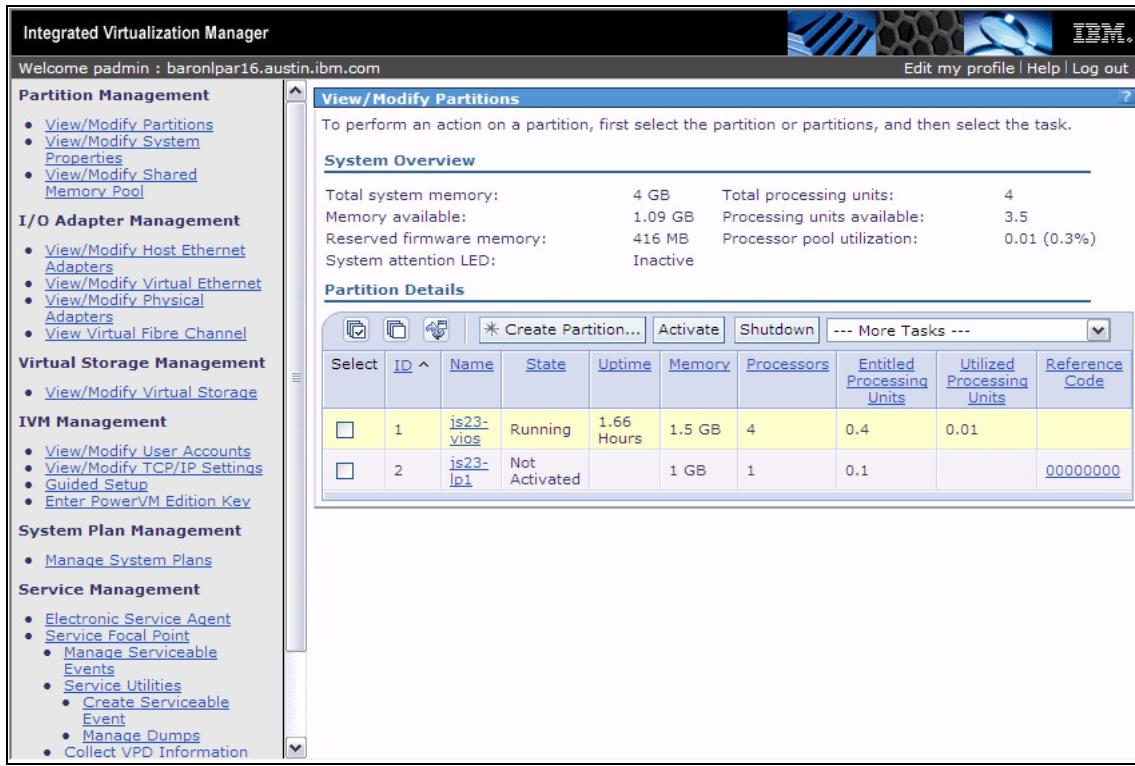


Figure 4-6 IVM navigation and work areas

The login to the UI is described in 4.4.1, “Connecting to IVM” on page 93

4.2.3 VIOS/IVM command line interface

The command line interface (CLI) requires more experience to master than the GUI, but it offers more possibilities for tuning the partition’s definitions. It can also be automated through the use of scripts.

The text-based console with the CLI is accessible through a Serial over LAN (SoL) terminal connection to the blade through the Advanced Management Module (AMM) or through network connectivity using Telnet or Secure Shell (SSH).

The login to the CLI is to a restricted shell. The restricted shell can be identified by the \$ prompt. A list of available commands can be displayed by using the `help` command. Individual VIOS command help can be viewed by using the `-help` flag

and IVM command help by using the **--help** flag. Detailed command help can be shown using the **man** command.

Note: Not all IVM commands will be displayed using the **help** command. For a complete listing of these commands, refer to *Virtual I/O Server and Integrated Virtualization Manager Command Reference*, which is available from:

<http://publib.boulder.ibm.com/infocenter/systems/topic/iphcg/iphcg.pdf>

4.3 First VIOS login

During the first login to the VIOS running on a IBM BladeCenter JS23 or JS43, you will be prompted to complete tasks that are required to make the server operational and IVM accessible through a Web browser. This first connection can be made from a Serial Over Lan (SoL) or KVM session.

4.3.1 Password set

From the login prompt, use the default user ID of **padmin**. You will be required to enter a password and then re-enter the same password for confirmation, as shown in Example 4-1.

4.3.2 License acceptance

The licensing agreements must be accepted prior to starting any configuration tasks. After setting the password, you will be presented three license options as shown in Example 4-1. Enter **v** to view the license, **d** to decline, or **a** to accept.

Example 4-1 Initial login screen showing password set and licensing options

IBM Virtual I/O Server

```
login: padmin
[compat]: 3004-610 You are required to change your password.
Please choose a new one.
```

```
padmin's New password:
Enter the new password again:
```

Indicate by selecting the appropriate response below whether you accept or decline the software maintenance terms and conditions.

After you enter a, enter the **license -accept** command as shown in Example 4-2.

Example 4-2 The license command

```
$ license -accept
```

The status of the license can be verified by using the **license** command with no flags, as shown in Example 4-3.

Example 4-3 The license status

```
$ license
The license has been accepted
en_US Apr 2 2009, 12:33:16 10(padmin)
```

4.3.3 Initial network setup

IVM requires a valid network configuration to be accessed by a Web browser. The VIOS IP address can be set by two different methods: by using the **mktcpip** command, or by starting a SMIT-like tool by using the **cfgassist** command.

Note: If you are reinstalling VIOS on the same blade and want to remove the previous logical partition information and return the blade to an original unconfigured state prior to making any configuration changes, execute the following command:

```
lpcfgop -o clear
```

The LPAR information will now be cleared on the next system restart.

The help function in the IVM CLI does not display the **lpcfgop** command. However, you can obtain the command's description by using:

```
man lpcfgop
```

Use the **lscppip -adapters** command to determine the available network adapters, as shown in Example 4-4 on page 87. When configuring an interface be sure to pick an adapter that has physical access to the network such as a Logical Host Ethernet Port and not a Virtual I/O Ethernet Adapter.

Example 4-4 The lscppip -adapters command used to display available network adapters

```
$ lscppip -adapters
```

```
Ethernet adapters:  
ent0 Available Logical Host Ethernet Port (1p-hea)  
ent1 Available Logical Host Ethernet Port (1p-hea)  
ent2 Available Virtual I/O Ethernet Adapter (1-lan)  
ent3 Available Virtual I/O Ethernet Adapter (1-lan)  
ent4 Available Virtual I/O Ethernet Adapter (1-lan)  
ent5 Available Virtual I/O Ethernet Adapter (1-lan)  
ent6 Available 05-20 Gigabit Ethernet-SX PCI-X Adapter (14106703)  
ent7 Available 05-21 Gigabit Ethernet-SX PCI-X Adapter (14106703)  
ibmvmc0 Available Virtual Management Channel  
$
```

Choose the corresponding interface on an adapter (ent0 and en0, in this example) that will be the target of the TCP/IP configuration.

Using the **mktcpip** command method

The **mktcpip** command has the following syntax:

```
mktcpip -hostname HostName -inetaddr Address -interface Interface  
[-start] [-netmask SubnetMask] [-cabletype CableType]  
[-gateway Gateway] [-nsrvaddr NameServerAddress  
-nsrvdomain Domain] [-plen prefixLength]
```

The **-cabletype** and **-plen** flags are the only two optional parameters. Example 4-5 shows how to configure a VIOS network interface.

Example 4-5 The mktcpip command

```
$ mktcpip -hostname saturn -inetaddr 172.16.1.200 -interface en0 -start  
-netmask 255.255.255.0 -gateway 172.16.1.1 -nsrvaddr 172.16.1.199  
-nsrvdomain customer.com
```

Using the **cfgassist** menu method

The second method using the **cfgassist** tool requires the same network information, but is entered in a more user-friendly method.

Example 4-6 on page 88 shows the initial **cfgassist** menu. You start the tool by entering the **cfgassist** command.

Example 4-6 The cfgassist menu

```
Config Assist for VIOS
```

Move cursor to desired item and press Enter.

- Set Date and TimeZone
- Change Passwords
- Set System Security
- VIOS TCP/IP Configuration
- Install and Update Software
- Storage Management
- Devices
- Electronic Service Agent

Esc+1=Help **Esc+2=Refresh** **Esc+3=Cancel** **F8=Image**
F9=Shell **F10=Exit** **Enter=Do**

By selecting **VIOS TCP/IP Configuration**, you will be presented a list of available network interfaces as shown in Example 4-7.

Example 4-7 cfgassist VIOS Available Network Interfaces

Config Assist for VIOS

Select the desired interface. On the next screen, shown in Example 4-8, you enter the TCP/IP configuration by pressing the Enter key. This completes the initial TCP/IP configuration of the VIOS.

Example 4-8 cfgassist TCP/IP interface configuration entry page

VIOS TCP/IP Configuration

Type or select values in entry fields.

Press Enter AFTER making all desired changes.

[Entry Fields]	
* Hostname	[saturn]
* Internet ADDRESS (dotted decimal)	[172.16.1.200]
Network MASK (dotted decimal)	[255.255.255.0]
* Network INTERFACE	en0
Default Gateway (dotted decimal)	[172.16.1.1]
NAMESERVER	
Internet ADDRESS (dotted decimal)	[172.16.1.199]
DOMAIN Name	[customer.com]
CableType	tp
Esc+1=Help	Esc+2=Refresh
Esc+5=Reset	F6=Command
F9=Shell	F10=Exit
	Esc+3=Cancel
	F7>Edit
	Enter=Do
	Esc+4>List
	F8=Image

Reviewing the Network Configuration

After you configure a network adapter, you can review the settings by using the **1stcpip** command. The **1stcpip** command has the following syntax:

```
Usage: 1stcpip [-num] [ [-routtable] | [-routinfo] | [-state] | [-arp]
]
          [-sockets [-family inet | inet6 | unix] ]
          [-stored] [-adapters] [-hostname] [-namesrv]
      1stcpip [-state [-field FieldName ...] ] |
          [-routtable [-field FieldName ...] ] [-fmt delimiter]
[-num]
      1stcpip [-namesrv] | [-interfaces] [-fmt delimiter]
```

Example 4-9 on page 91 shows sample output for the **-adapters**, **-interfaces**, **-routtable**, and **-stored** flags as they are used with the **1stcpip** command to show basic TCP/IP configuration.

Example 4-9 lscplip command sample output

```
$ lscplip -interfaces
```

Name	Address	Netmask	State	MAC
en1	-	-	detach	00:1a:64:76:00:09
en3	-	-	detach	2a:08:4b:c1:cb:04
en4	-	-	detach	2a:08:4b:c1:cb:05
en5	-	-	detach	2a:08:4b:c1:cb:06
et1	-	-	detach	00:1a:64:76:00:09
et3	-	-	detach	2a:08:4b:c1:cb:04
et4	-	-	detach	2a:08:4b:c1:cb:05
et5	-	-	detach	2a:08:4b:c1:cb:06
en6	172.16.1.200	255.255.255.0	up	00:1a:64:76:00:08
et6	-	-	detach	00:1a:64:76:00:08

```
$ lscplip -routtable
```

Routing tables

Destination	Gateway	Flags	Refs	Use	If	Exp
-------------	---------	-------	------	-----	----	-----

Groups

Route Tree for Protocol Family 2 (Internet):

default	172.16.1.1	UG	2	288871	en6	-
-						
127/8	localhost	U	9	209	lo0	-
-						
172.16.1.0	saturn.customer.com	UHSb	0	0	en6	
- - - =>						
172.16.1/24	saturn.customer.com	U	2	5628596	en6	
- - -						
saturn.ibm.com	localhost	UGHS	0	39074	lo0	-
-						
172.16.1.255	saturn.customer.com	UHSb	0	4	en6	
- - -						

Route Tree for Protocol Family 24 (Internet v6):

::1	::1	UH	0	232	lo0	-
-						

```
$ lscplip -stored
```

saturn

Network Interface Attributes

attributes: en0

```
IPv4 address = 172.16.1.200
Network Mask = 255.255.255.0
State = detach

attributes: en1
State = down

attributes: en2
State = down

attributes: en3
State = down

attributes: en4
State = down

attributes: en5
State = down

attributes: et0
State = detach

attributes: et1
State = down

attributes: et2
State = down

attributes: et3
State = down

attributes: et4
State = down

attributes: et5
State = down

attributes: en6
IPv4 address = 172.16.1.200
Network Mask = 255.255.255.0
State = up

attributes: et6
State = detach
```

```
attributes: en7
    State = down

attributes: en8
    State = down

attributes: et7
    State = down

attributes: et8
    State = down

Static Routes:
Route 1:
    hopcount = 0
    default gateway = 172.16.1.1

DNS information:
    nameserver 172.16.1.199
    domain customer.com
```

To remove all or part of the TCP/IP configuration, use the **rmtcpip** command. The **rmtcpip** command has the following syntax:

```
Usage: rmtcpip [-f] [-nextboot] {-all | [-hostname] [-routing]
[-interface ifnameList]}
rmtcpip [-f] {-all | [-namesrv] [-hostname] [-routing]
[-interface ifnameList]}
```

To remove all TCP/IP configuration, use the **rmtcpip -f -all** command.

4.4 First IVM connection

When the initial networking tasks are completed, the IVM UI should be accessible through a Web browser. The VIOS Web Server supports HTTP and HTTPS connections.

4.4.1 Connecting to IVM

Using a Web browser window, connect using HTTP or HTTPS to the IP address that you assigned to the VIOS during the installation process (see 4.3.3, “Initial network setup” on page 87).

A Welcome window that contains the login and password prompts opens, as shown in Figure 4-7. The default user ID is padmin, and the password is the one you defined during the VIOS installation.

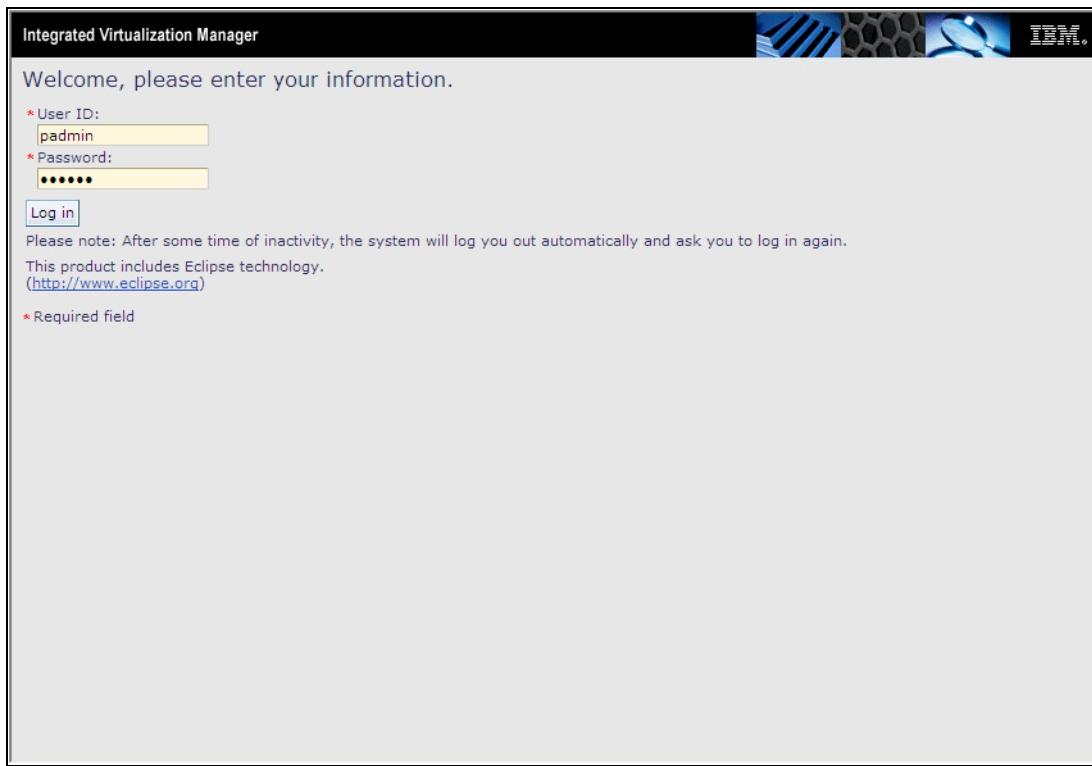


Figure 4-7 The Welcome window

The first connection to the IVM UI will display the guided setup window as shown in Figure 4-8 on page 95. Expanding the sections on the window provides additional information about configuration and management tasks, with links directly to some of the functions. You can return to this window at any time by clicking the **Guided Setup** link in the navigation area.

Before configuring any additional LPARs or resources, review the default installation values of the VIOS.

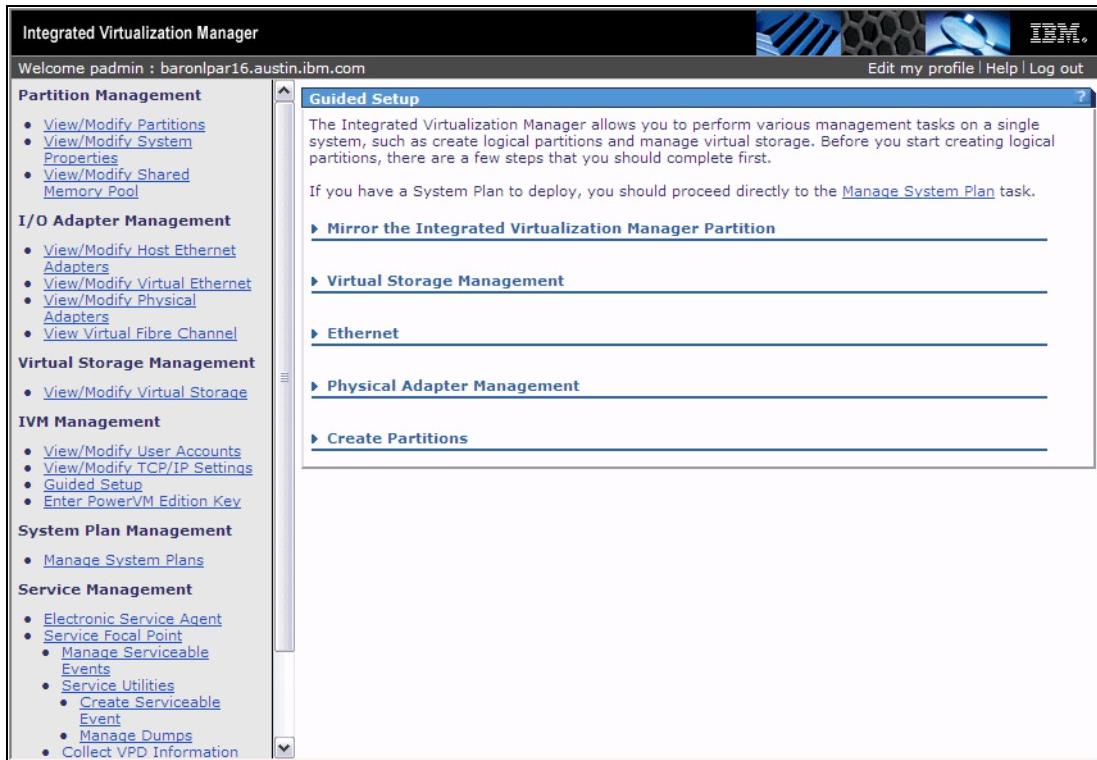


Figure 4-8 Guided Setup window

4.4.2 Verify and modify VIOS partition memory and processors

After the initial installation of the VIOS, there is only one LPAR, the VIOS, on the system with the following characteristics:

- ▶ The ID is 1.
- ▶ The name is equal to the system's serial number.
- ▶ The state is Running.
- ▶ The allocated memory is between 1GB and one-eighth of the installed system memory.
- ▶ The number of (virtual) processors is equal to the number of cores, and the Entitled Processing Units is equal to 0.1 times the number of allocated processors.

Details of all configured partitions, including the VIOS, are summarized on the View/Modify Partitions window as shown in Figure 4-9. This window can be accessed by the **View/Modify Partitions** link in the navigation area.

The screenshot shows the Integrated Virtualization Manager (IVM) interface. On the left, a sidebar lists various management categories with their sub-links. The main area is titled "View/Modify Partitions" and contains sections for "System Overview" and "Partition Details".

System Overview:

Total system memory:	4 GB	Total processing units:	4
Memory available:	2.62 GB	Processing units available:	3.6
Reserved firmware memory:	384 MB	Processor pool utilization:	0.01 (0.3%)
System attention LED:	Inactive		

Partition Details:

Select	ID	Name	State	Uptime	Memory	Processors	Entitled Processing Units	Utilized Processing Units	Reference Code
<input type="checkbox"/>	1	10-1821A	Running	3.65 Days	1 GB	4	0.4	0.01	

Figure 4-9 View/Modify Partitions window

Administrators can change properties of the VIOS LPAR, including memory or processing units allocation by using the IVM UI. From the View/Modify Partitions window, click the link in the Name column that corresponds to ID 1 (The VIOS will always be ID or LPAR 1).

The Partition Properties window will be displayed in a new window, as shown in Figure 4-10 on page 97. The name of the VIOS can be changed from the General tab, if desired. The Memory and Processing tabs are used to view or change the allocations. Figure 4-11 on page 98 shows the General tab.

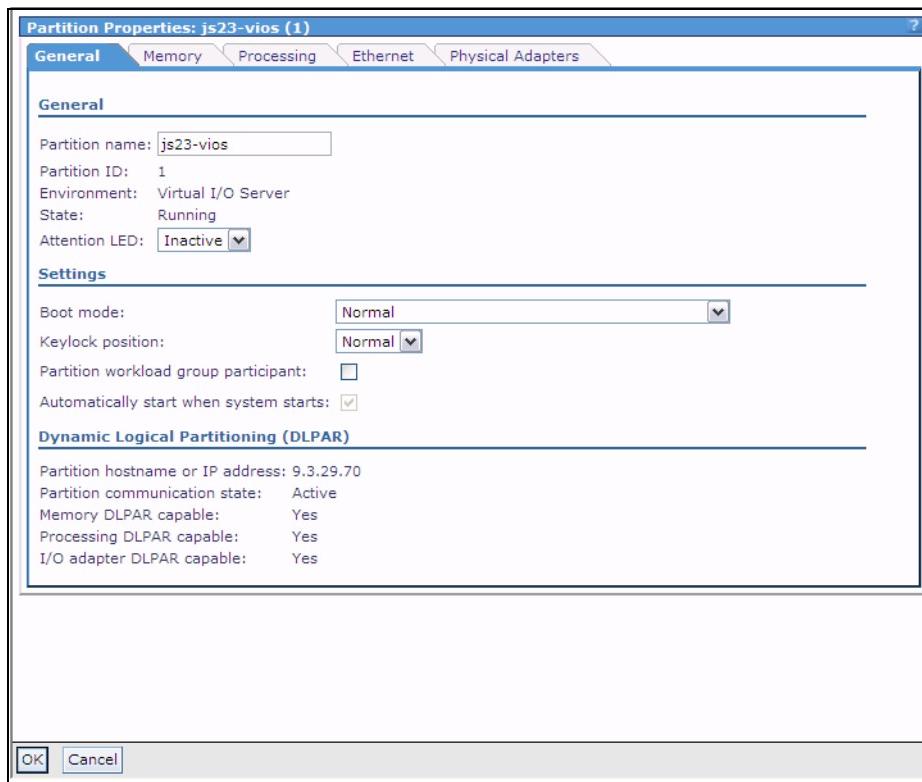


Figure 4-10 Partition Properties, General tab

Figure 4-11 shows the Memory tab.

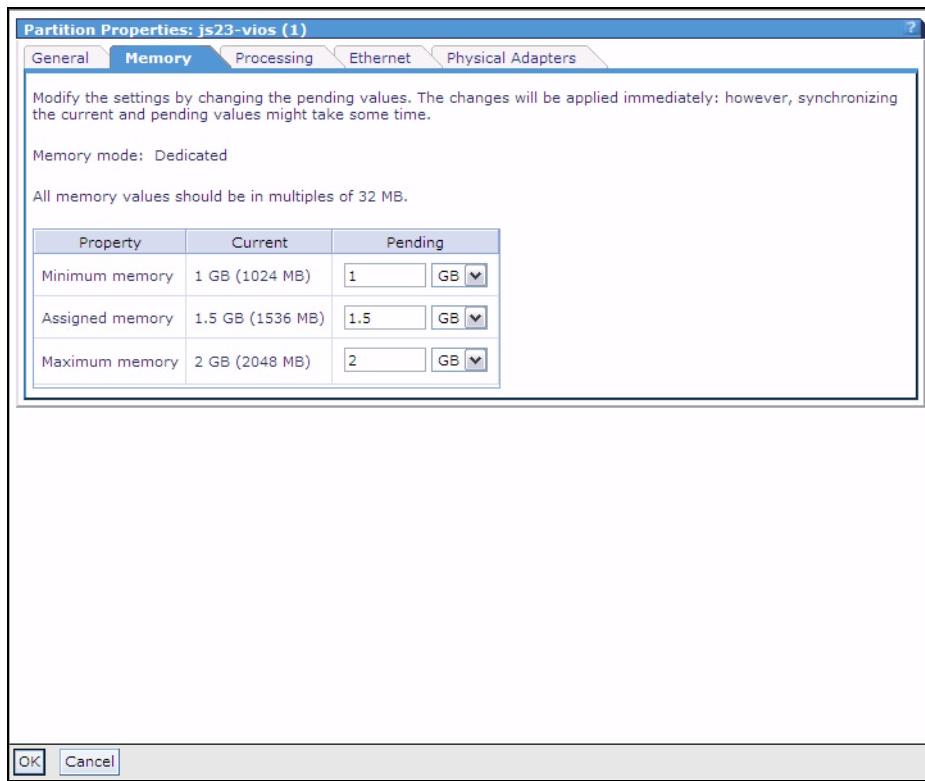


Figure 4-11 Partition Properties, Memory tab

The default memory configuration for the VIOS LPAR, 1/8 of system memory with a minimum value of 1GB. You may need to increase memory values if it did default to 1GB and you are using additional expansion cards or combinations of expansion cards and EtherChannel configurations, or you plan to have an LPAR supporting IBM i partition. The Assigned memory value should not be reduced below the default minimum of 1GB.

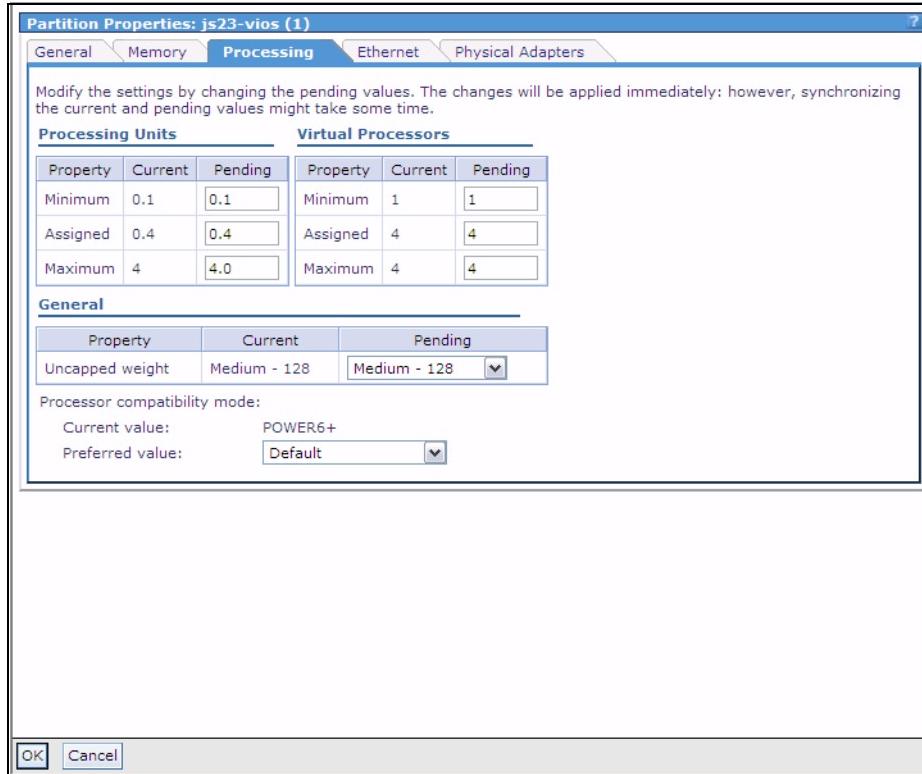


Figure 4-12 Partition Properties, Processing tab

Processing unit allocations for the VIOS are recommended to remain at the install defaults. But you should monitor utilization and adjust the Assigned amount, as required. The Virtual Processor default settings should not be changed.

The **lshwres** and **chsyscfg** commands are used to display and change memory and processor values, as shown in Example 4-10.

Example 4-10 CLI commands for displaying and altering memory and processor values

```
$ lshwres -r mem --level lpar --filter "\"lpar_ids=1\"" -F curr_mem
1024

$ chsyscfg -r prof -i "lpar_id=1,desired_mem=1152"

$ lshwres -r mem --level lpar --filter "\"lpar_ids=1\"" -F curr_mem
1152
```

```
$ lshwres -r proc --level lpar --filter "\"lpar_ids=1\"" -F  
curr_proc_units  
0.40  
  
$ chsyscfg -r prof -i "lpar_id=1,desired_proc_units=0.5"  
  
$ lshwres -r proc --level lpar --filter "\"lpar_ids=1\"" -F  
curr_proc_units  
0.50  
  
$ lshwres -r proc --level lpar --filter "\"lpar_ids=1\"" -F curr_procs  
4  
$ chsyscfg -r prof -i "lpar_id=1,desired_procs=3"  
$ lshwres -r proc --level lpar --filter "\"lpar_ids=1\"" -F curr_procs  
3
```

The Ethernet tabs are discussed in 4.5.2, “Virtual Ethernet Adapters and SEA” on page 103. Physical Adapters tabs are discussed in 4.5.3, “Physical adapters” on page 109.

4.5 VIOS network management and setup

When you are preparing the IBM BladeCenter JS23 or JS43 and VIOS for additional LPARs, and planning for the ability for those LPARs to reach a physical network, you need to understand the networking options that are available. The three choices are:

- ▶ Logical Host Ethernet Adapters (Ip-HEA)
- ▶ Virtual Ethernet Adapters bridged to a physical adapters via a Shared Ethernet Adapter (SEA)
- ▶ Physical Adapters

4.5.1 Host Ethernet Adapters

Using the Host Ethernet Adapters (HEA) will reduce overall system CPU usage by the VIOS compare to a SEA, but will not allow partition mobility. In a JS23 16 individual logical ports in one port group are available across the two physical ports. The JS43 has 32 logical ports over four physical ports and two ports groups. In either case the VIOS will claim one logical port on each physical port leaving 14 available logical ports on the JS23 for LPAR assignment and 28 logical ports on the JS43. The assignment of a logical port from the HEA can be

done during the LPAR creation process. Refer to 3.8.4, “Integrated Virtual Ethernet (IVE)” on page 56 for additional technical details about the HEA.

You configure the HEA port mode by selecting **View/Modify Host Ethernet Adapters** from the navigation area. This displays the UI window, as shown in Figure 4-13 on page 101.

The screenshot shows the IVM interface with the following details:

- Header:** Integrated Virtualization Manager, Welcome padmin : baronlpar47.austin.ibm.com, IBM logo, Edit my profile | Help | Log out.
- Navigation Area (Left):**
 - Partition Management: View/Modify Partitions, View/Modify System Properties, View/Modify Shared Memory Pool.
 - I/O Adapter Management: View/Modify Host Ethernet Adapters, View/Modify Virtual Ethernet Adapters, View/Modify Physical Adapters, View Virtual Fibre Channel.
 - Virtual Storage Management: View/Modify Virtual Storage.
 - IVM Management: View/Modify User Accounts, View/Modify TCP/IP Settings, Guided Setup, Enter PowerVM Edition Key.
 - System Plan Management: Manage System Plans.
 - Service Management: Electronic Service Agent, Service Focal Point (with sub-options: Manage Serviceable Events, Service Utilities (Create Serviceable Event, Manage Dumps), Collect VPD Information).
- Central Window:** Title: View/Modify Host Ethernet Adapters. Subtitle: A Host Ethernet Adapter (HEA) allows you to provide multiple partitions direct access to the physical Ethernet ports. To perform an action on a physical port, first select the port, and then select the task. A table titled "Properties" lists four ports:

Select	Type	Link State	Physical Location	Code	Connected Partitions	Available Connections
<input type="radio"/>	1 G	Up	U78A5.001.WIH23CF-P1-T6	1		0
<input type="radio"/>	1 G	Up	U78A5.001.WIH23CF-P1-T7	2		13
<input type="radio"/>	1 G	Up	U78A5.001.WIH23CF-P2-T6	1		14
<input type="radio"/>	1 G	Up	U78A5.001.WIH23CF-P2-T7	1		0

Figure 4-13 View/Modify Host Ethernet Adapters window

All four HEA ports on a JS43 are shown. The default configuration is port sharing with 14 logical connections available per port pair. In this example two port have been set to allow bridging (Available Connections 0) one port has an additional partition connected (Connected Partitions 2).

To view and modify the port properties, select one of the ports and click **Properties**. The HEA Physical Port Properties window will be displayed in a new window; Figure 4-14 on page 102 shows the General tab already selected.

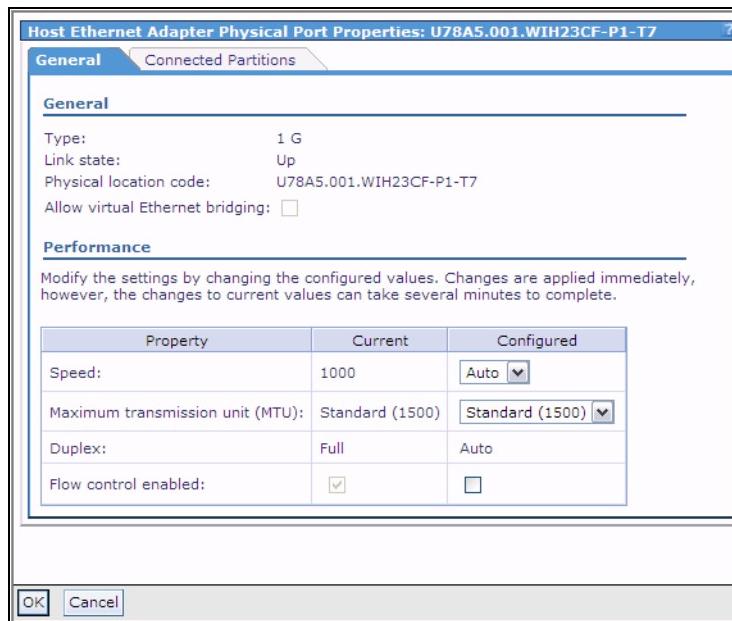


Figure 4-14 HEA Port Properties

You can display a list of connected partitions (if any) and MAC addresses by selecting the **Connected Partitions** tab, as shown in Figure 4-15 on page 103.

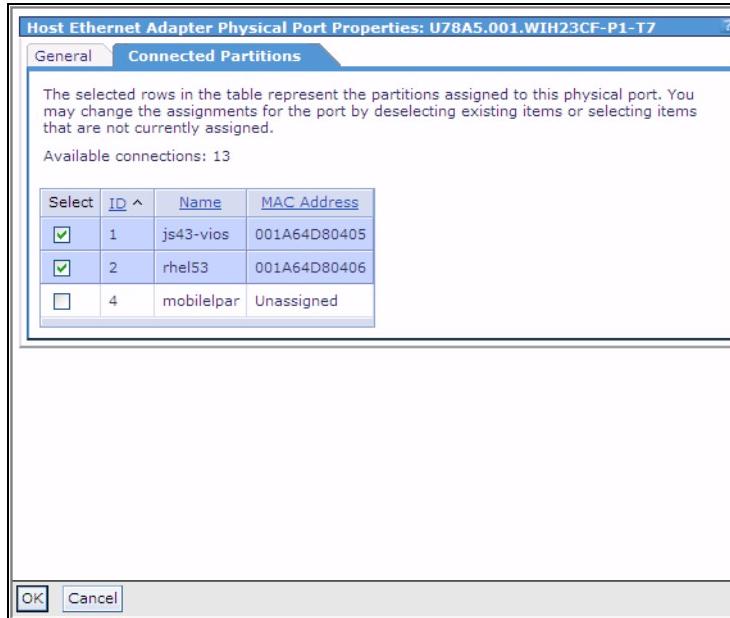


Figure 4-15 HEA Port Properties, Connected Partitions

4.5.2 Virtual Ethernet Adapters and SEA

Virtual adapters exist in the hypervisor that allows LPARs to communicate with each other without the need for a physical network. They can be created for each partition provided by the hypervisor.

Four virtual Ethernet adapters are created by default on the VIOS, and two each for every logical partition. Additional virtual adapters can be created on both the VIOS and logical partitions.

A Shared Ethernet Adapter (SEA) is a new virtual adapter is created by bridging between a physical adapter (HEA port or expansion card port) and a virtual adapter on the VIOS. A SEA adapter can also be created by bridging between an EtherChannel adapter and a virtual adapter. The SEA interface can then be mapped to logical partitions, thus providing network connectivity outside of the VIOS and logical partition.

For a HEA adapter port to participate in an SEA Ethernet bridge, the configuration of the port must be changed. The configuration is changed by ensuring that the box **Allow virtual Ethernet bridging** on the HEA Physical Port Properties window is checked, as shown in Figure 4-16 on page 104. This setting

allows the port to operate in promiscuous mode. When this mode is enabled, there is only one logical port available and it is assigned to the VIOS LPAR.

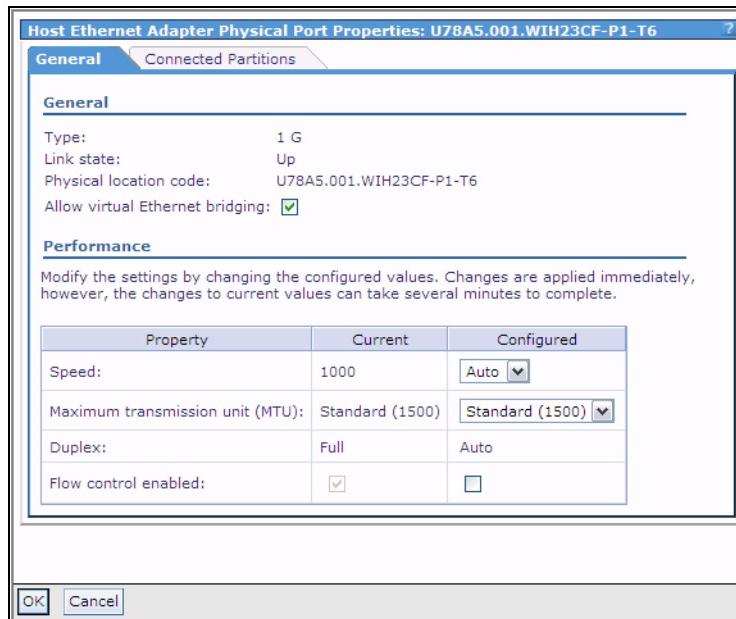


Figure 4-16 HEA port setting for Ethernet bridging

Physical Ethernet ports on an expansion cards do not require configuration prior to being used in a SEA environment.

The SEA adapter is configured by selecting the **View/Modify Virtual Ethernet** link in the navigation area. If the four default virtual Ethernet adapter have not been already created you will see the screen as shown in Figure 4-17 on page 105. Click the **Initialize Virtual Ethernet** button to create the default adapters.

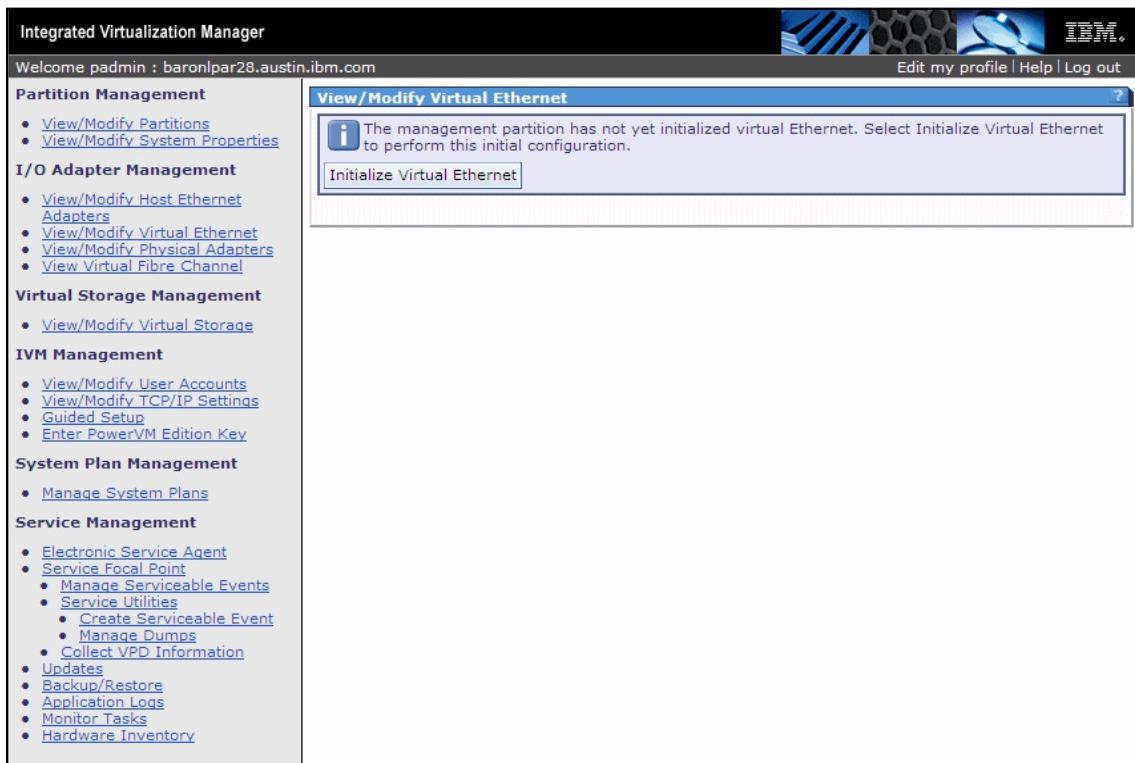


Figure 4-17 View/Modify Virtual Ethernet showing Initialize Virtual Ethernet option

Figure 4-18 on page 106 shows the four virtual Ethernet adapters that are created by default on the VIOS.

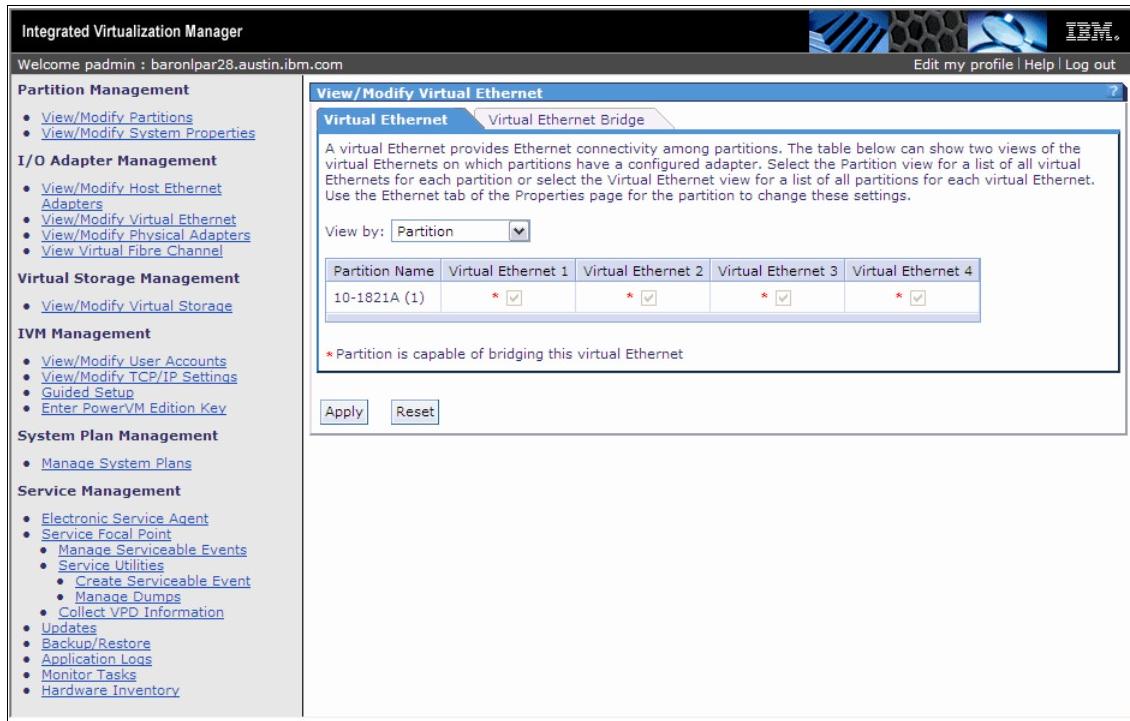


Figure 4-18 View/Modify Virtual Ethernet window

Use the **Virtual Ethernet Bridge** tab to display the virtual to physical options for creating an SEA, as shown in Figure 4-19 on page 107. The drop-down box in the Physical Adapter column lists the adapters that are available for creating the SEA.

Notes:

- ▶ A physical adapter can only be used to create one SEA in combination with a virtual adapter. The drop-down box in the UI will allow the selection of the same adapter for another virtual Ethernet ID, but an error message will be generated when the **Apply** button is clicked.
- ▶ EtherChannel adapters, if created, will also be listed in the drop-down box.

Integrated Virtualization Manager

Welcome padmin : baronpar28.austin.ibm.com

Partition Management

- [View/Modify Partitions](#)
- [View/Modify System Properties](#)

I/O Adapter Management

- [View/Modify Host Ethernet Adapters](#)
- [View/Modify Virtual Ethernet](#)
- [View/Modify Physical Adapters](#)
- [View Virtual Fibre Channel](#)

Virtual Storage Management

- [View/Modify Virtual Storage](#)

IVM Management

- [View/Modify User Accounts](#)
- [View/Modify TCP/IP Settings](#)
- [Guided Setup](#)
- [Enter PowerVM Edition Key](#)

System Plan Management

- [Manage System Plans](#)

Service Management

- [Electronic Service Agent](#)
- [Service Focal Point](#)
 - [Manage Serviceable Events](#)
 - [Service Utilities](#)
 - [Create Serviceable Event](#)
 - [Manage Dumps](#)
 - [Collect VPD Information](#)
 - [Updates](#)
 - [Backup/Restore](#)
 - [Application Logs](#)
 - [Monitor Tasks](#)
 - [Hardware Inventory](#)

View/Modify Virtual Ethernet

Virtual Ethernet **Virtual Ethernet Bridge**

A physical network bridge provides a specific virtual Ethernet access to a physical Ethernet device, thereby allowing any partition on the bridged virtual Ethernet to access the external network via the physical Ethernet device. For a given virtual Ethernet, you may choose the physical adapter to which to bridge. If a virtual Ethernet is not bridged, traffic on the virtual Ethernet is restricted to member partitions.

Virtual Ethernet ID	Physical Adapter
1	None
2	None
3	None
4	None

Apply **Reset**

Figure 4-19 View/Modify Virtual Ethernet Bridge tab

Figure 4-20 on page 108 shows a physical adapter selection.

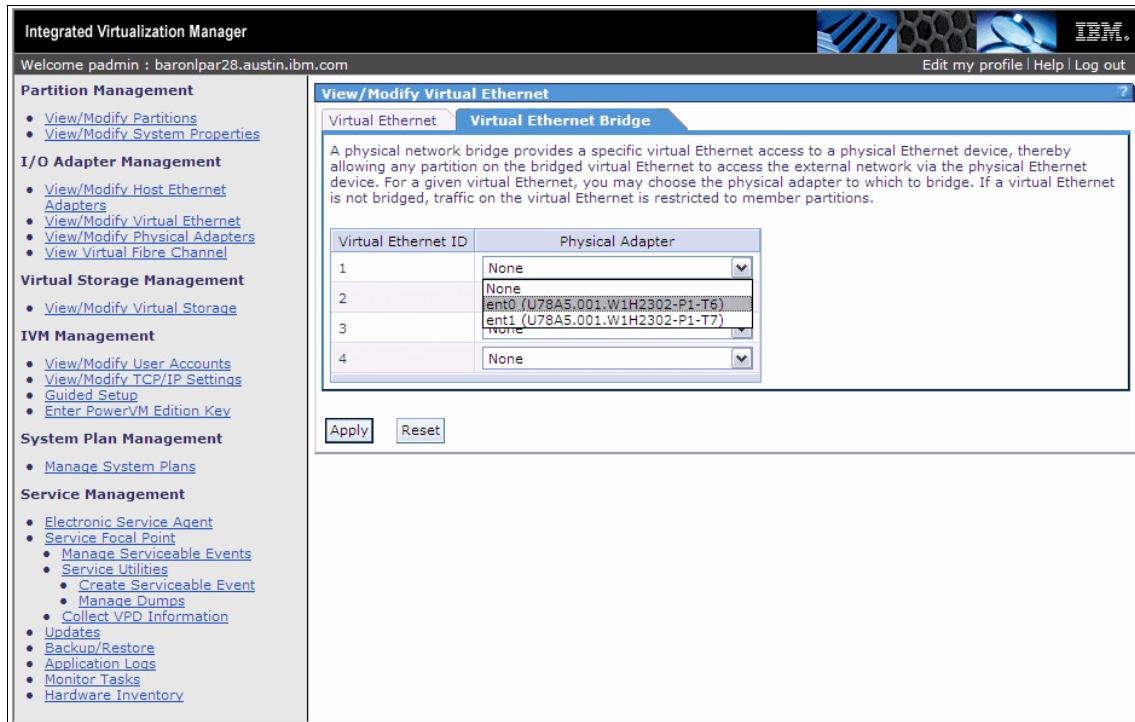


Figure 4-20 Physical adapter selection for SEA creation

Figure 4-21 on page 109 indicates the successful creation of the SEA.

The screenshot shows the IVM interface with the following details:

- Left Navigation Bar:**
 - Integrated Virtualization Manager**
 - Welcome padmin : baronpar28.austin.ibm.com
 - Partition Management**
 - [View/Modify Partitions](#)
 - [View/Modify System Properties](#)
 - I/O Adapter Management**
 - [View/Modify Host Ethernet Adapters](#)
 - [View/Modify Virtual Ethernet](#)
 - [View/Modify Physical Adapters](#)
 - [View Virtual Fibre Channel](#)
 - Virtual Storage Management**
 - [View/Modify Virtual Storage](#)
 - IVM Management**
 - [View/Modify User Accounts](#)
 - [View/Modify TCP/IP Settings](#)
 - [Guided Setup](#)
 - [Enter PowerVM Edition Key](#)
 - System Plan Management**
 - [Manage System Plans](#)
 - Service Management**
 - [Electronic Service Agent](#)
 - [Service Focal Point](#)
 - [Manage Serviceable Events](#)
 - [Service Utilities](#)
 - [Create Serviceable Event](#)
 - [Manage Dumps](#)
 - [Collect VPD Information](#)
 - [Updates](#)
 - [Backup/Restore](#)
 - [Application Logs](#)
 - [Monitor Tasks](#)
 - [Hardware Inventory](#)
- Main Content Area:**

View/Modify Virtual Ethernet

Virtual Ethernet **Virtual Ethernet Bridge**

Information Message: The operation completed successfully.

A physical network bridge provides a specific virtual Ethernet access to a physical Ethernet device, thereby allowing any partition on the bridged virtual Ethernet to access the external network via the physical Ethernet device. For a given virtual Ethernet, you may choose the physical adapter to which to bridge. If a virtual Ethernet is not bridged, traffic on the virtual Ethernet is restricted to member partitions.

Virtual Ethernet ID	Physical Adapter
1	ent0 (U78A5.001.W1H2302-P1-T6)
2	None
3	None
4	None

Buttons: Apply | Reset

Figure 4-21 Successful SEA creation result

4.5.3 Physical adapters

With the IBM BladeCenter JS23 or JS43, you have the option to assign physical hardware adapters to an LPAR. From a network perspective, only Ethernet expansion cards can be reassigned to an LPAR. The HEA adapter ports cannot be assigned to a logical partition.

Note: When using IBM i and shared memory partitions, the resources must be purely virtual.

To assign a physical adapter, select the **View/Modify Physical Adapters** link from the navigation area to display the window shown in Figure 4-22 on page 110.

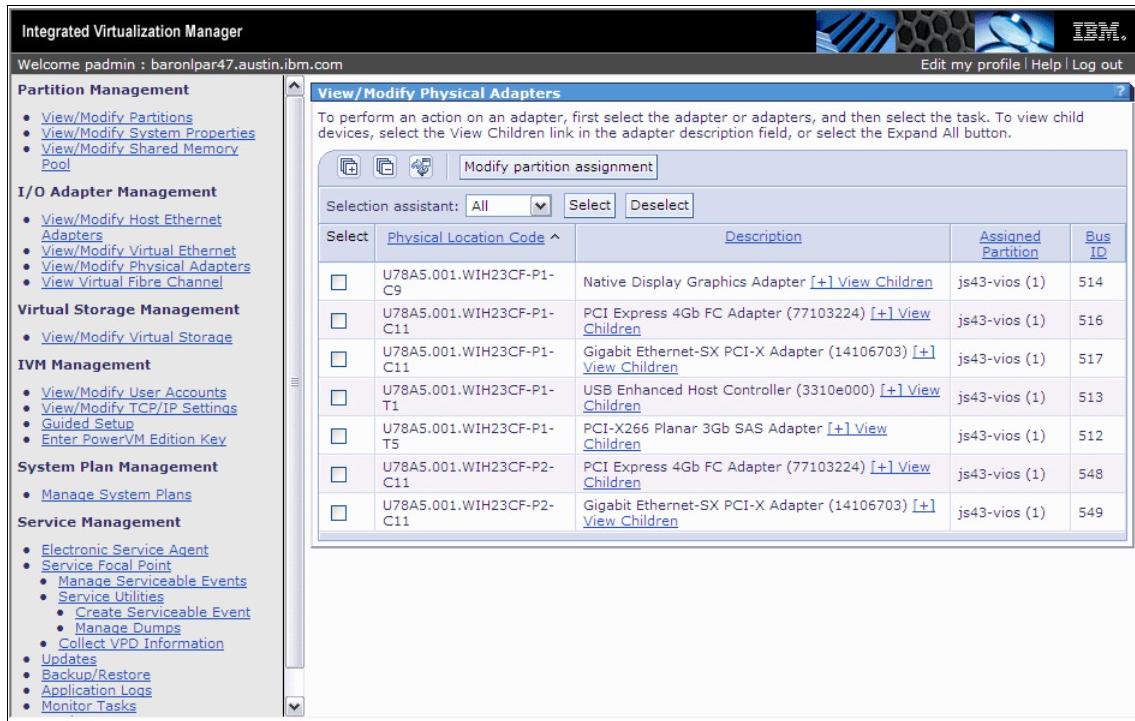


Figure 4-22 View/Modify Physical Adapters window

By default, all physical adapters are owned by the VIOS LPAR. By using the **Modify Partition Assignment** button, you can change the assigned partition.

In the example shown in Figure 4-23 on page 111, the Gigabit Ethernet expansion card ports are being reassigned to partition 2.

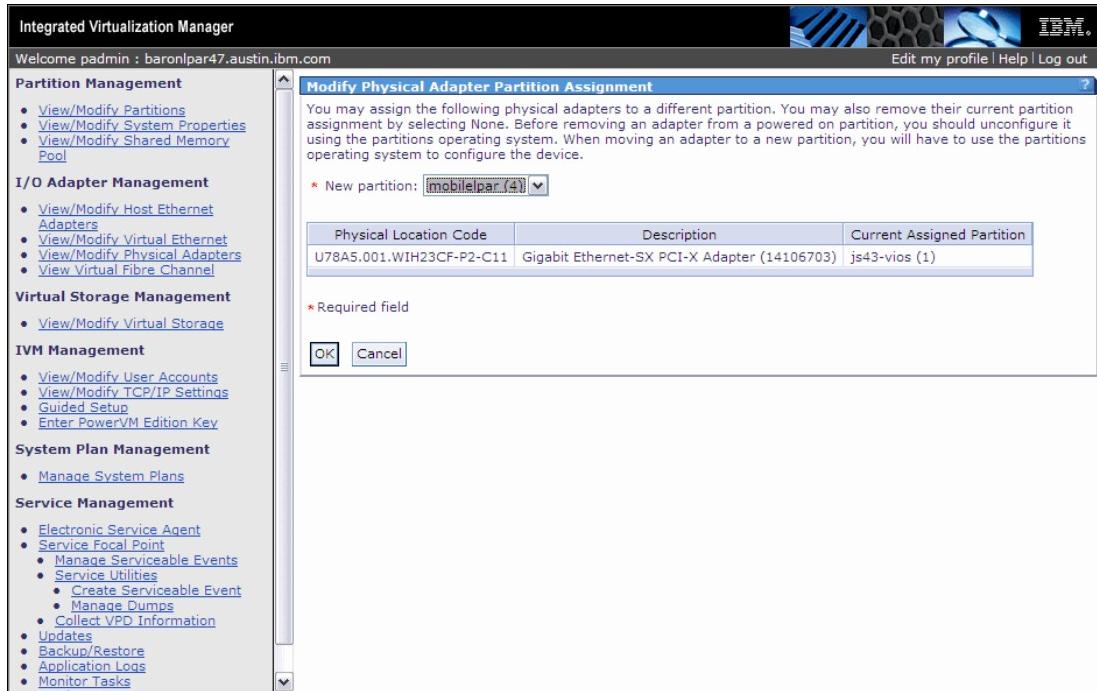


Figure 4-23 Physical Adapter assignment to new partition

Figure 4-24 on page 112 shows the change in partition ownership.

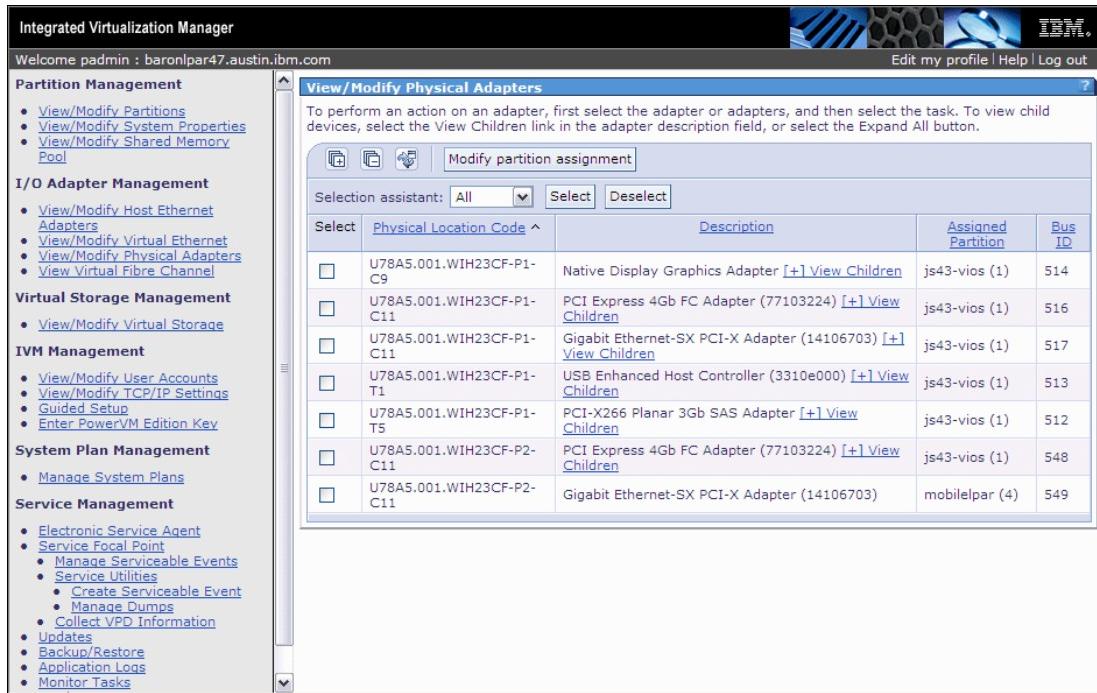


Figure 4-24 View/Modify Physical Adapter window showing change of ownership of Gigabit Ethernet Adapter

Example 4-11 shows the changes in adapter availability in an AIX logical partition, starting with the original virtual Ethernet adapter through the addition of the two physical ports from an IBM BladeCenter JS23 or JS43 expansion card.

Example 4-11 Physical port addition to logical partition

```
# lsdev -Cc adapter
ent0 Available Virtual I/O Ethernet Adapter (1-lan)
vsa0 Available LPAR Virtual Serial Adapter
vsclsi0 Available Virtual SCSI Client Adapter
# lsdev -Cc adapter
ent0 Available Virtual I/O Ethernet Adapter (1-lan)
ent1 Defined 01-20 Gigabit Ethernet-SX PCI-X Adapter (14106703)
ent2 Defined 01-21 Gigabit Ethernet-SX PCI-X Adapter (14106703)
vsa0 Available LPAR Virtual Serial Adapter
vsclsi0 Available Virtual SCSI Client Adapter
# cfgmgr
# lsdev -Cc adapter
ent0 Available Virtual I/O Ethernet Adapter (1-lan)
ent1 Available 01-20 Gigabit Ethernet-SX PCI-X Adapter (14106703)
```

```
ent2 Available 01-21 Gigabit Ethernet-SX PCI-X Adapter (14106703)
vsa0 Available LPAR Virtual Serial Adapter
vsccsi0 Available Virtual SCSI Client Adapter
#
```

Note: When removing a physical adapter from a LPAR you may have to remove a PCI bus device with a `rmdev` command from the LPAR's command line. The IVM interface will display an error message with text indicating the device that must be removed before the change in LPAR assignment can be performed.

4.5.4 VLANs and virtual networks

When the VIO Server is installed four virtual Ethernet adapters are created automatically. These adapters have Port VLAN Identifiers (PVID) of 1-4. These PVIDs are considered untagged and do not leave the VIO Server. In order to extend a VLAN environment into the virtual network. The concept of VLAN tagging using IEEE 802.1q protocol is used to make the extension.

IVM does not provide a direct way to create 802.1q aware interfaces but can manage the assignment of these interfaces after they are created using the CLI.

The steps to perform, and the VIO Server user interfaces used are:

- ▶ Create IEEE 802.1q aware virtual Ethernet adapter on the VIOS that will act as a trunking adapter to carry multiple VLANs (CLI)
- ▶ Create a SEA between the physical adapter and the VLAN aware virtual Ethernet adapter on the VIO Server (IVM)
- ▶ Assign the LPAR virtual Ethernet adapter to the VIO Server virtual Ethernet
- ▶ Run `cfgmgr` on LPAR (if running) and configure new interfaces (client LPAR CLI)

Section 3.6 in *IBM System p Advanced POWER Virtualization (PowerVM) Best Practices*, REDP-4194 contains additional discussions on extending VLANs into virtual networks. The implementation covered in this Redpaper is for a Hardware Management Console or HMC managed system not IVM, but the technical discussions of VLANs in a virtual network are relevant.

Another example of implementation from an IBM BladeCenter and VIOS/IVM perspective, *Complex networking using Linux on Power blades*, can be found at this link:

<http://www.ibm.com/developerworks/power/library/l-bladenetconf/index.html?ca=drs->

VLAN configuration of BladeCenter Ethernet switch modules or other Ethernet switches external to the BladeCenter are not covered in this document.

Creating new VIOS virtual Ethernet adapters

The four default virtual adapters that are created by the VIO Server during installation cannot be modified for VLAN tagging use. Therefore new virtual adapters must be created using the CLI with the desire VLAN information.

The **lsdev** command can be used to review the current real and virtual Ethernet adapters on our lab VIO Server as shown Example 4-12 on page 114.

Example 4-12 lsdev command used to list current Ethernet adapters

```
$ lsdev -type adapter |grep ent
ent0          Available  Logical Host Ethernet Port (1p-hea)
ent1          Available  Logical Host Ethernet Port (1p-hea)
ent2          Available  Virtual I/O Ethernet Adapter (1-lan)
ent3          Available  Virtual I/O Ethernet Adapter (1-lan)
ent4          Available  Virtual I/O Ethernet Adapter (1-lan)
ent5          Available  Virtual I/O Ethernet Adapter (1-lan)
ent6          Available  Gigabit Ethernet-SX PCI-X Adapter
ent7          Available  Gigabit Ethernet-SX PCI-X Adapter
ent8          Available  Shared Ethernet Adapter
```

Using the **lshwres** command we can view the existing virtual Ethernet resources. In Example 4-13 the four default virtual Ethernet adapters are shown. The output gives details of the adapter characteristics such as LPAR ownership, PVIDs (1-4), additional VLANs (none), and IEEE 802.1q capability (0=no).

Example 4-13 lshwres command showing VIO Server virtual Ethernet adapters

```
$ lshwres -r virtualio --rsubtype eth --level lpar
lpar_name=js23-vios,lpar_id=1,slot_num=3,state=1,ieee_virtual_eth=0,por
t_vlan_id=1,addl_vlan_ids=none,is_trunk=1,trunk_priority=1,is_required=
0,mac_addr=067E5E2D8C03
lpar_name=js23-vios,lpar_id=1,slot_num=4,state=1,ieee_virtual_eth=0,por
t_vlan_id=2,addl_vlan_ids=none,is_trunk=1,trunk_priority=1,is_required=
0,mac_addr=067E5E2D8C04
lpar_name=js23-vios,lpar_id=1,slot_num=5,state=1,ieee_virtual_eth=0,por
t_vlan_id=3,addl_vlan_ids=none,is_trunk=1,trunk_priority=1,is_required=
0,mac_addr=067E5E2D8C05
```

```
1par_name=js23-vios,1par_id=1,slot_num=6,state=1,ieee_virtual_eth=0,por  
t_vlan_id=4,addl_vlan_ids=none,is_trunk=1,trunk_priority=1,is_required=  
0,mac_addr=067E5E2D8C06
```

From the **View/Modify Virtual Ethernet** view in IVM as shown Figure 4-25 the four default VIO Server Ethernet adapters are displayed.

Partition Name	Virtual Ethernet 1	Virtual Ethernet 2	Virtual Ethernet 3	Virtual Ethernet 4
js23-vios (1)	* <input checked="" type="checkbox"/>			
IBMI 2 (2)	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
VLANpar3 (3)	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>

Figure 4-25 Default VIO Server virtual Ethernet Adapters shown by IVM

Note: Figure 4-25 shows additional partitions. Partition creation is not covered until 4.7, “Partition configuration for Virtual I/O Client (VIOC)” on page 144.

When creating a new virtual Ethernet adapter for the VIO Server we will be assigning PVIDs and additional VLAN ids. The PVID used should be unique and not used by any clients in the network or physical Ethernet switch ports. This requirement is to prevent the unintentional removal or stripping of VLAN tags from network packets when entering the VIO Server.

The **chhwres** command is used to created a new VIO Server virtual Ethernet adapter. Example 4-14 creates a new VIOS adapter with a PVID of 555, with IEEE 802.1q is enabled and additional VLANS of 20, 30, and 40,

Example 4-14 Using chhwres command to create new VIOS virtual Ethernet adapter

```
$ chhwres -r virtualio --rsubtype eth -o a --id 1 -s 15 -a  
port_vlan_id=555,ieee_virtual_eth=1,\"add1_vlan_ids=20,30,40\",is_trunk  
=1,trunk_priority=1
```

The flags and their attributes are:

- ▶ **-r virtualio --rsubtype eth** type of hardware resource to change
- ▶ **-o a** perform add operation
- ▶ **--id 1** the LPAR id number
- ▶ **-s 15** slot number to use
- ▶ **-a** attributes to add
 - **port_vlan_id=555** PVID
 - **ieee_virtual_eth=1** turns on IEEE 802.1q support
 - **add_vlan_ids=20,30,40** defines additional VLAN ids
 - **is_trunk=1** must be turned on to pass multiple VLANs
 - **trunk_priority=1** priority of the adapter can be set between 1-15

The **lsdev** command is repeated and shows the new Virtual I/O Ethernet Adapter at ent9 as shown in Example 4-15

Example 4-15 sdev command used to list new ent9 Ethernet adapters

```
$ lsdev -type adapter |grep ent  
ent0      Available  Logical Host Ethernet Port (1p-hea)  
ent1      Available  Logical Host Ethernet Port (1p-hea)  
ent2      Available  Virtual I/O Ethernet Adapter (1-lan)  
ent3      Available  Virtual I/O Ethernet Adapter (1-lan)  
ent4      Available  Virtual I/O Ethernet Adapter (1-lan)  
ent5      Available  Virtual I/O Ethernet Adapter (1-lan)  
ent6      Available  Gigabit Ethernet-SX PCI-X Adapter  
ent7      Available  Gigabit Ethernet-SX PCI-X Adapter  
ent8      Available  Shared Ethernet Adapter  
ent9      Available  Virtual I/O Ethernet Adapter (1-lan)
```

The **lshwres** command is run again as shown in Example 4-16 to display the new adapter indicated in bold.

Example 4-16 lshwres command showing new VIOS virtual Ethernet adapter

```
$ lshwres -r virtualio --rsubtype eth --level lpar
```

```

lpar_name=js23-vios,lpar_id=1,slot_num=3,state=1,ieee_virtual_eth=0,por
t_vlan_id=1,addl_vlan_ids=none,is_trunk=1,trunk_priority=1,is_required=
0,mac_addr=067E5E2D8C03
lpar_name=js23-vios,lpar_id=1,slot_num=4,state=1,ieee_virtual_eth=0,por
t_vlan_id=2,addl_vlan_ids=none,is_trunk=1,trunk_priority=1,is_required=
0,mac_addr=067E5E2D8C04
lpar_name=js23-vios,lpar_id=1,slot_num=5,state=1,ieee_virtual_eth=0,por
t_vlan_id=3,addl_vlan_ids=none,is_trunk=1,trunk_priority=1,is_required=
0,mac_addr=067E5E2D8C05
lpar_name=js23-vios,lpar_id=1,slot_num=6,state=1,ieee_virtual_eth=0,por
t_vlan_id=4,addl_vlan_ids=none,is_trunk=1,trunk_priority=1,is_required=
0,mac_addr=067E5E2D8C06
lpar_name=js23-vios,lpar_id=1,slot_num=15,state=1,ieee_virtual_eth=1,po
rt_vlan_id=555,"addl_vlan_ids=20,30,40",is_trunk=1,trunk_priority=1,is_
required=0,mac_addr=067E5E2D8C0F

```

With a new VIO Server virtual Ethernet adapter created, we are ready to use IVM for the next steps in configuration.

Create SEA using IEEE802.1q virtual Ethernet adapter

IVM will now display the new virtual Ethernet adapter in the **View/Modify Virtual Ethernet** view as shown in Figure 4-26 on page 117. Notice the PVID and additional VLAN numbers are shown to be associated with this new virtual adapter.

Partition Name	Virtual Ethernet 1	Virtual Ethernet 2	Virtual Ethernet 3	Virtual Ethernet 4	Other
js23-vios (1)	* <input checked="" type="checkbox"/>	* 20, * 30, * 40, * 555			
IBMI 2 (2)	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	
VLANpar3 (3)	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	

* Partition is capable of bridging this virtual Ethernet

Figure 4-26 View Modify Virtual Ethernet view showing new virtual adapter with multiple VLANs

We now create a SEA or bridge between this new virtual adapter and a physical Ethernet port, in this case a HEA adapter, by first clicking the **Virtual Ethernet Bridge** tab. From the virtual Ethernet list we choose 555(20,30,40) and map it to ent1 as shown in Figure 4-27 on page 118. Click **OK** to complete the assignment and the creation of the SEA.

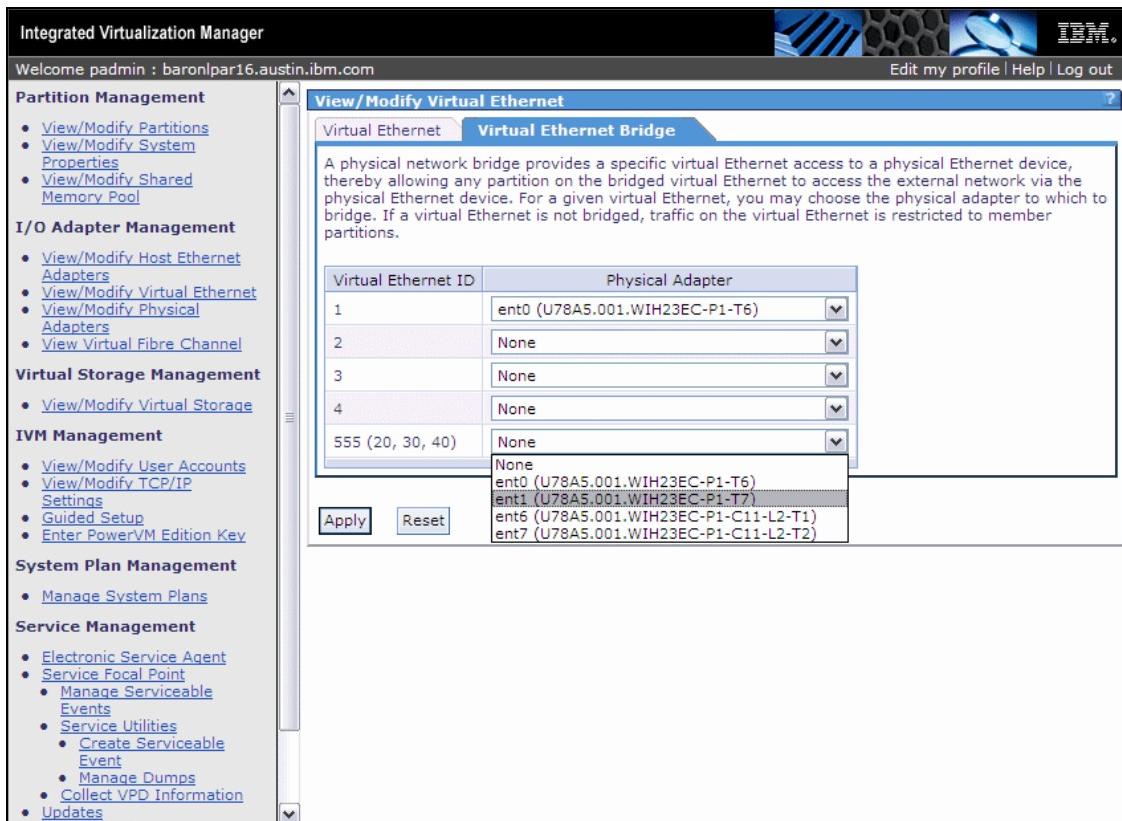


Figure 4-27 Creating a SEA using an IEEE 802.1q enabled virtual adapter

Using the **lsdev** command we see the new VIO Server virtual Ethernet adapter ent9, and the new SEA ent10 as shown in Example 4-17

Example 4-17 lsdev command showing new IEEE 802.1q virtual adapter and SEA

```
$ lsdev -type adapter |grep ent
ent0          Available  Logical Host Ethernet Port (lp-hea)
ent1          Available  Logical Host Ethernet Port (lp-hea)
ent2          Available  Virtual I/O Ethernet Adapter (l-lan)
ent3          Available  Virtual I/O Ethernet Adapter (l-lan)
ent4          Available  Virtual I/O Ethernet Adapter (l-lan)
```

ent5	Available	Virtual I/O Ethernet Adapter (1-1an)
ent6	Available	Gigabit Ethernet-SX PCI-X Adapter
ent7	Available	Gigabit Ethernet-SX PCI-X Adapter
ent8	Available	Shared Ethernet Adapter
ent9	Available	Virtual I/O Ethernet Adapter (1-1an)
ent10	Available	Shared Ethernet Adapter

With the successful creation of the SEA we can use the **entstat** command on the VIO Server to get additional details of the components of the SEA as shown in Example 4-18.

Example 4-18 entstat command used to provide VLAN details

```
$ entstat -all ent10 |grep VLAN
VLAN Ids :
    VLAN Extract: False
    VLAN tagged filtering mode: Filter according to VLAN permit array
Max number of VLAN IDs per HEA port: 20
Invalid VLAN ID Packets: 0
Port VLAN ID: 555
VLAN Tag IDs: 40    30    20
```

VIO Client LPAR virtual Ethernet mapping

The next step is to map a VIO Client LPAR virtual Ethernet adapter to the VIO Server virtual Ethernet. This procedure is done using IVM. Click the LPAR name from the **View/Modify Partitions** view to bring up the Partition Properties window then click the **Ethernet** tab. As shown in Figure 4-28 on page 120 in the Virtual Ethernet Adapter section, pick a VIOC adapter, and from the corresponding pull down box pick the desired VLAN id. In this example we used VIOC adapter 2 and VIOS virtual Ethernet corresponding to VLAN 20. When the selection is complete click the **OK** button.

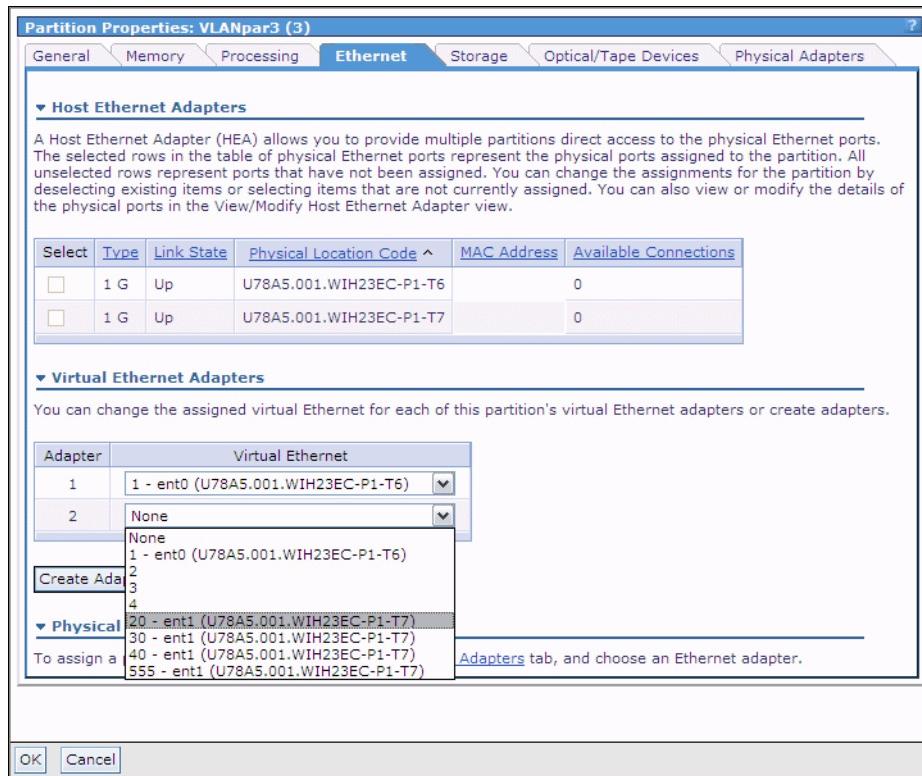


Figure 4-28 VIOC adapter to VIOS virtual Ethernet mapping

VIO Client verification and configuration

If the partition is not active, the new adapter will be discovered upon activation of the LPAR. If the partition is already active you may need to take additional steps such as run the `cfgmgr` command in AIX. IBM i LPARs with Autoconfig enabled will automatically configure the new adapter.

With the discovery of the new virtual adapter for the LPAR complete a new virtual Ethernet adapter, ent1, is available in our lab example VIO Client as shown in Example 4-19.

Example 4-19 lsdev command from VIO Client showing new virtual Ethernet

```
# lsdev |grep ent
ent0      Available      Virtual I/O Ethernet Adapter (1-lan)
ent1      Available      Virtual I/O Ethernet Adapter (1-lan)
```

VLAN details of ent1 can be displayed using the **entstat** command on the VIO Client (assumes an AIX client) as shown in Example 4-20 on page 121.

Example 4-20 entstat command from VIO Client showing details of new virtual Ethernet

```
# entstat -d ent1 |grep VLAN
Invalid VLAN ID Packets: 0
Port VLAN ID:    20
VLAN Tag IDs:  None
```

In this AIX LPAR example The interface en1 on VLAN 20 can now be configured with the desired TCP/IP properties

4.6 VIOS Storage Configuration and Management

Virtual Storage Management allows you to manage the disk storage requirements for logical partitions. The following terms are used to describe storage management concepts:

- ▶ Physical volumes
- ▶ Storage pools
- ▶ Virtual disks
- ▶ Optical and Tape devices

These terms are explained in the next sections.

To work with VIOS storage, click **View/Modify Virtual Storage** in the navigation area of the IVM as shown in Figure 4-29.

System Overview

Total system memory:	8 GB	Total processing units:	8
Memory available:	3.28 GB	Processing units available:	5.3
Reserved firmware memory:	640 MB	Processor pool utilization:	0.07 (0.9%)
System attention LED:	Inactive		

Partition Details

Select	ID ^	Name	State	Uptime	Memory	Processors	Entitled Processing Units	Utilized Processing Units	Reference Code
<input type="checkbox"/>	1	js43-vios	Running	1.95 Days	1.5 GB	8	0.8	0.02	
<input type="checkbox"/>	2	rhel53	Running	7.4 Hours	1 GB	1	0.1	0.03	Linux ppc64
<input type="checkbox"/>	3	IBMi	Running	2.16 Days	1 GB	1	1.0	0.00	00000000
<input type="checkbox"/>	4	mobilepar	Running	1.92 Days	608 MB	8	0.8	0.02	

Figure 4-29 View and modify virtual storage

4.6.1 Physical volumes

Physical volumes are the hard drives that are available to the VIOS. They can be installed locally in the IBM BladeCenter JS23 or JS43 blades, SAS drives available from IBM BladeCenter S chassis, or LUNs available from a Fibre Channel storage area network subsystem.

A physical volume is shown as hdisk0, hdisk1 and so on. The Virtual I/O Server. Logical partitions can be assigned complete physical volumes. However, they will appear as a virtual SCSI disk drive on the LPAR. This direct assignment is a requirement if you are planning for Live Partition Mobility, shared memory partitions or using IBM i. Physical volumes can also be used to build storage pools or AMS dedicated paging devices.

To verify the available physical volumes in your VIOS, in the navigation area, click **View/Modify Virtual Storage**. Then click the **Physical Volumes** tab, as shown

in Figure 4-30 on page 123. This displays the list of the physical volumes available to the VIOS.

Select	Name	Storage Pool	Assigned Partition	Size	Physical Location Code
<input type="checkbox"/>	hdisk0	rootvg (Default)		68.37 GB	U78A5.001.WIH23CF-P1-T5-L20000-L0
<input type="checkbox"/>	hdisk1			68.37 GB	U78A5.001.WIH23CF-P1-T5-L60000-L0
<input type="checkbox"/>	hdisk2	media_pool		30 GB	U78A5.001.WIH23CF-P1-C11-L1-T1-W500507630E87FE3F-L40104001000000
<input type="checkbox"/>	hdisk3			30 GB	U78A5.001.WIH23CF-P1-C11-L1-T1-W500507630E87FE3F-L40104002000000
<input type="checkbox"/>	hdisk4			30 GB	U78A5.001.WIH23CF-P1-C11-L1-T1-W500507630E87FE3F-L40104003000000
<input type="checkbox"/>	hdisk5			30 GB	U78A5.001.WIH23CF-P1-C11-L1-T2-W500507630E85FE3F-L40104004000000
<input type="checkbox"/>	hdisk6			20 GB	U78A5.001.WIH23CF-P1-C11-L1-T2-W500507630E85FE3F-L40104011000000
<input type="checkbox"/>	hdisk7			20 GB	U78A5.001.WIH23CF-P1-C11-L1-T1-W500507630E87FE3F-L40104012000000
<input type="checkbox"/>	hdisk8	IBMi (3)		20 GB	U78A5.001.WIH23CF-P1-C11-L1-T1-W500507630E87FE3F-L40104013000000
<input type="checkbox"/>	hdisk9	IBMi (3)		20 GB	U78A5.001.WIH23CF-P1-C11-L1-T1-W500507630E87FE3F-L40104014000000
<input type="checkbox"/>	hdisk10	mobilelpar (4)		15 GB	U78A5.001.WIH23CF-P1-C11-L1-T1-W500507630E87FE3F-L4010401F000000
<input type="checkbox"/>	hdisk11			15 GB	U78A5.001.WIH23CF-P1-C11-L1-T1-W500507630E87FE3F-L4010402000000
<input type="checkbox"/>	hdisk12			15 GB	U78A5.001.WIH23CF-P1-C11-L1-T1-W500507630E87FE3F-L40104021000000
<input type="checkbox"/>	hdisk13			15 GB	U78A5.001.WIH23CF-P1-C11-L1-T1-W500507630E87FE3F-L40104022000000
<input type="checkbox"/>	hdisk14	rhel53 (2)		15 GB	U78A5.001.WIH23CF-P1-C11-L1-T1-W500507630E87FE3F-L40104023000000

Figure 4-30 Physical volumes shown in IVM

Similar information can be retrieved on the Virtual I/O Server CLI by using the **lsdev** and **1spv** commands. Example 4-21 shows the output of the **lsdev -type disk** command.

Example 4-21 Physical volumes found with lsdev

```
$ lsdev -type disk
name          status    description
hdisk0        Available SAS Disk Drive
hdisk1        Available SAS Disk Drive
hdisk2        Available IBM MPIO FC 1750
hdisk3        Available IBM MPIO FC 1750
hdisk4        Available IBM MPIO FC 1750
hdisk5        Available IBM MPIO FC 1750
hdisk6        Available IBM MPIO FC 1750
hdisk7        Available IBM MPIO FC 1750
```

hdisk8	Available	IBM MPIO FC 1750
hdisk9	Available	IBM MPIO FC 1750
hdisk10	Available	IBM MPIO FC 1750
hdisk11	Available	IBM MPIO FC 1750
hdisk12	Available	IBM MPIO FC 1750
hdisk13	Available	IBM MPIO FC 1750
hdisk14	Available	IBM MPIO FC 1750

Example 4-22 shows the output of the `lspv -size` command.

Example 4-22 Physical volumes found with lspv -size

\$ lspv -size		
NAME	PVID	SIZE(megabytes)
hdisk0	000181ca0005e5c6	70006
hdisk1	000181ca6309a681	70006
hdisk2	000181ca7d20d77c	30720
hdisk3	none	30720
hdisk4	none	30720
hdisk5	none	30720
hdisk6	none	20480
hdisk7	none	20480
hdisk8	none	20480
hdisk9	none	20480
hdisk10	000180ea884b6253	15360
hdisk11	000180ea884b8500	15360
hdisk12	000180ea884b98b4	15360
hdisk13	000180ea884baacf	15360
hdisk14	000180ea884bbe1e	15360

4.6.2 Storage pools

A *storage pool* is a single entity that consists of one or more physical volumes or files. Logical volumes or virtual disks are created within a storage pool. Physical volumes or file spaces can only be assigned to a single storage pool.

After the installation of the VIOS, a default storage pool is created from space in the volume group `rootvg`. A new default storage pool should be created if you plan to use storage pools. As a general best practice, the storage pool in volume group `rootvg` should not be used.

Creating a new storage pool

To create a new storage pool, click the **Storage Pools** tab from the View/Modify Virtual Storage window. Figure 4-31 on page 125 shows a list of all available storage pools.

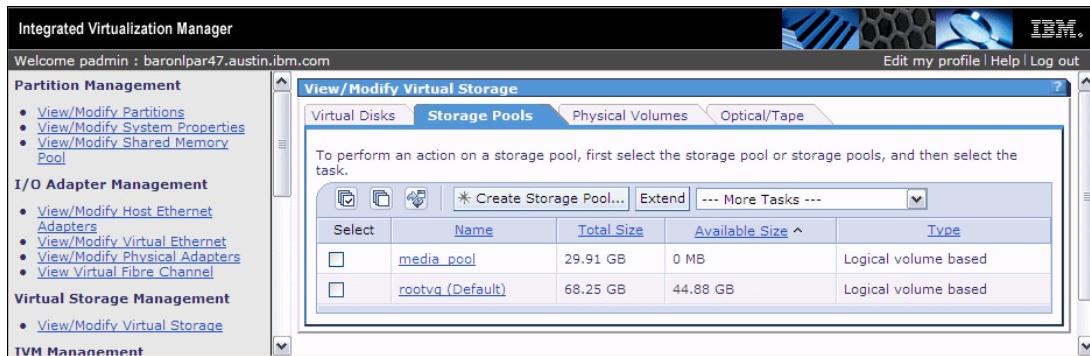


Figure 4-31 Storage pools shown in IVM

Click **Create Storage Pool...** to create a new storage pool. A dialog opens that guides you through the setup of the storage pool.

Specify a name (for example, SP-Media-Lib) that will be used for the storage pool. The name used for the storage pool must be a valid name for volume groups, for example no spaces are allowed are allowed and the name cannot exceed 15 characters.

Specify the storage pool type as Logical Volume based. The File based option currently requires a local file system. Select one or more available physical volumes to be used for the new storage pool, then click **OK**.

Figure 4-32 shows that, in this case, *hdisk3* was chosen.

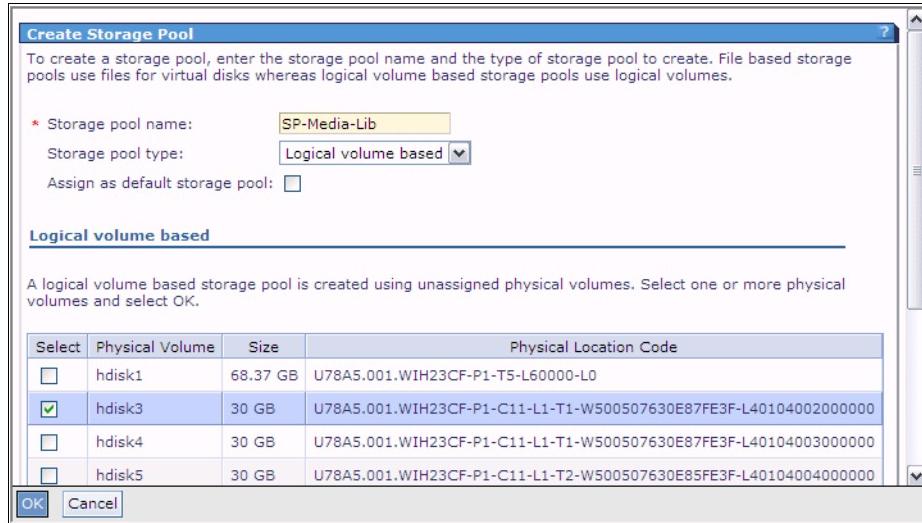


Figure 4-32 Create new storage pool

Figure 4-33 shows the new storage pool.

Select	Name	Total Size	Available Size	Type
<input type="checkbox"/>	media_pool	29.91 GB	0 MB	Logical volume based
<input type="checkbox"/>	rootvg (Default)	68.25 GB	44.88 GB	Logical volume based
<input type="checkbox"/>	SP-Media-Lib	29.91 GB	29.91 GB	Logical volume based

Figure 4-33 Newly created storage pool shown in IVM

Deleting or reducing a storage pool

To delete or reduce a storage pool, start from the Storage Pool tab in the Modify Virtual Storage window.

Select the storage pool you want to delete or reduce. Click **Reduce** from the More Tasks drop-down box as shown in Figure 4-34. A dialog opens that guides you through the modification of the storage pool.

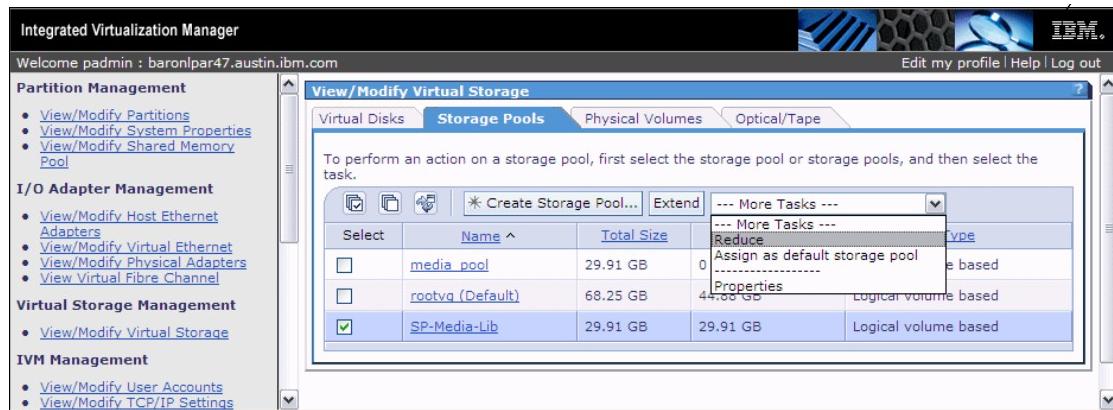


Figure 4-34 Reduce or delete a storage pool

Select the physical volumes that you want to remove from the storage pool. The storage pool will be deleted when all physical volumes that are assigned to the storage pool are removed. Click **OK**, as shown in Figure 4-35.

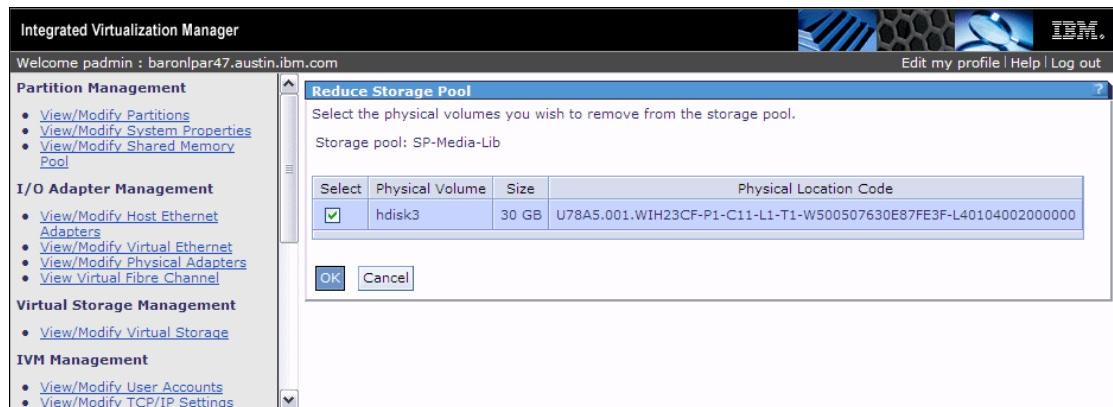


Figure 4-35 Delete storage pool

4.6.3 Virtual disks

Virtual disks are created in storage pools. After they are assigned to a logical partition, they are seen as virtual SCSI disk drives by the LPAR. These assignments are represented in the LPAR as hdisks. Multiple virtual disks can be created in a single storage pool. However, this method does not support Live Partition Mobility and is not recommended for IBM i.

You can create virtual disks from the View/Modify Virtual Storage window by selecting the **Virtual Disks** tab, as described in the following section. The Create Partition Wizard, as described in 4.7.2, “Partition name and environment” on page 145, can also be used to create virtual disks. Both methods require free space in a storage pool.

Creating virtual disks

To create a logical volume, a storage pool must be available. Refer to 4.6.2, “Storage pools” on page 124, for information about how to create a storage pool.

To create a new virtual disk, start with the View/Modify Virtual Storage window and select the **Virtual Disks** tab. From this window, click the **Create Virtual Disk** button as shown in Figure 4-36.

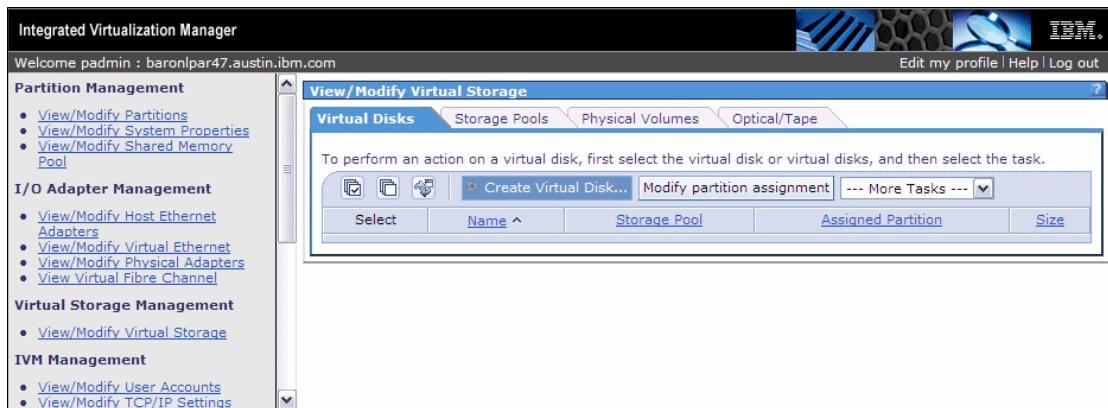


Figure 4-36 Create virtual disk

From the Create Virtual Disk window specify the name of the virtual disk, select a storage pool that will be used from the drop-down box, and specify the virtual disk size.

Optionally, you can make a partition assignment during virtual disk creation as shown in Figure 4-37 on page 129. A new virtual disk will be created when you click **OK**.

Note: When a virtual disk is created during the Create Partition wizard, the default naming schema for virtual disks uses the partition ID and the number of the assigned virtual disk to the LPAR. The result looks like `lp{number}vd{number}`.

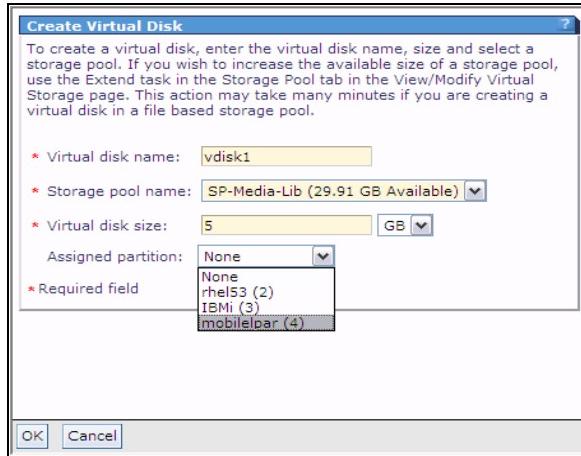


Figure 4-37 Virtual disk settings

The newly created virtual disk appears in the list, as shown in Figure 4-38.

Select	Name	Storage Pool	Assigned Partition	Size
<input type="checkbox"/>	vdisk1	SP-Media-Lib	mobilepar (4)	5 GB

Figure 4-38 The newly created virtual disk

The size of the virtual disk can be extended, as described in the following section.

Extending a virtual disk

You can extend a virtual disk as long as enough free space is available in the storage pool. To extend a virtual disk, select the virtual disk you plan to extend in the check box. Select the **More Tasks...** drop-down box, and then select **Extend**, as shown in Figure 4-39 on page 130.

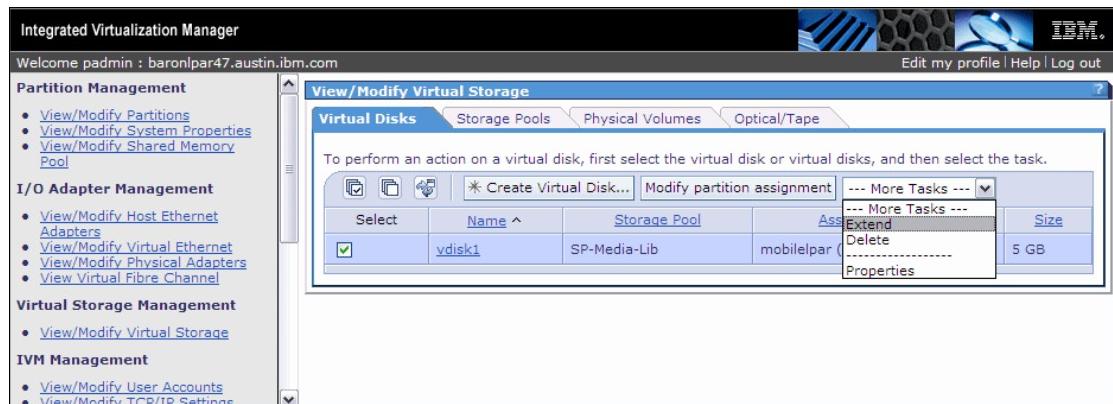


Figure 4-39 Extend virtual disk

Specify the amount of space that the virtual disk will be extended, then click **OK** as shown in Figure 4-40. If the storage pool does not have enough free space, it can be extended from the Storage Pools tab.

Note: When you attempt to extend virtual disk on a running partition, a warning message will be generated, alerting the administrator. To continue, select the **Force extend on running partition** check box and click the **OK** button again.

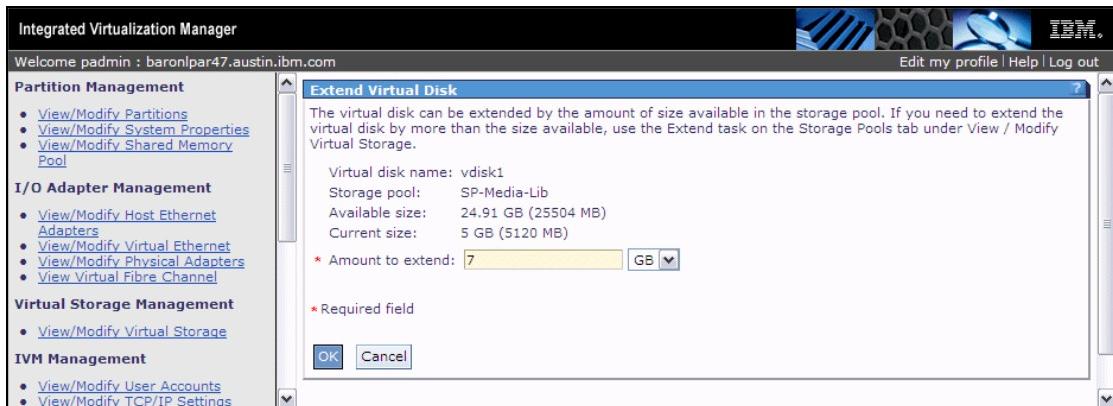


Figure 4-40 Virtual disk extension settings

The new size is shown in the list of available virtual disks when the extension is complete, as shown in Figure 4-41 on page 131.

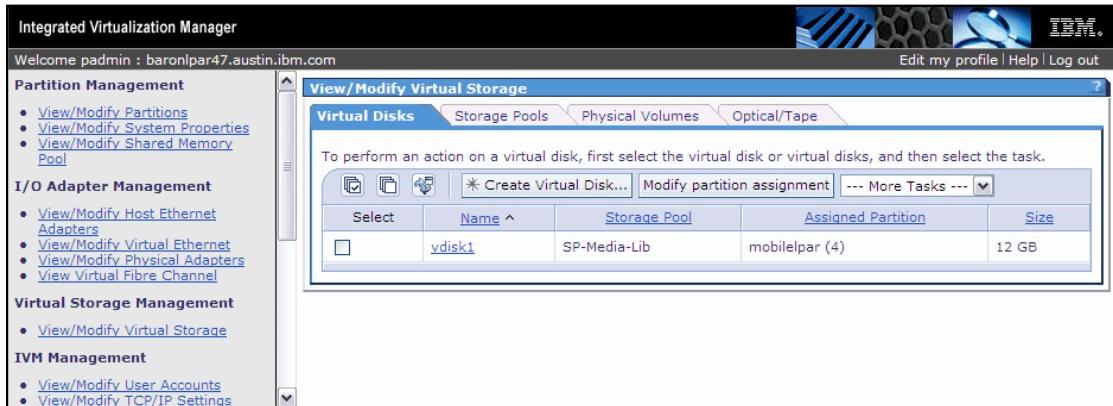


Figure 4-41 Extended virtual disk

The next section explains how to delete a virtual disk.

Deleting virtual disks

A virtual disk that is assigned to a partition must have that assignment removed before the virtual disk can be deleted.

Note: When you attempt to delete a virtual disk on a running partition, a warning message will be generated, alerting the administrator. To continue, select the **Force device removal from a running partition** check box and click the **OK** button again.

To delete a virtual disk click the **Virtual Disks** tab in the View/Modify Virtual Storage window. Select the virtual disk that you want to delete and select **Delete** from the **More Tasks...** drop down box as shown in Figure 4-42 on page 132.

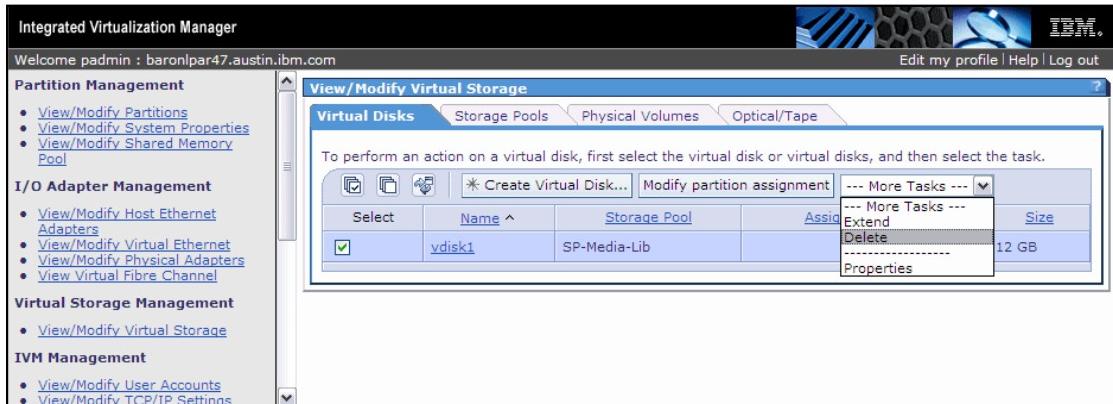


Figure 4-42 Delete virtual disk

Confirm the deletion of the virtual disk by clicking **OK**, as shown in Figure 4-43.

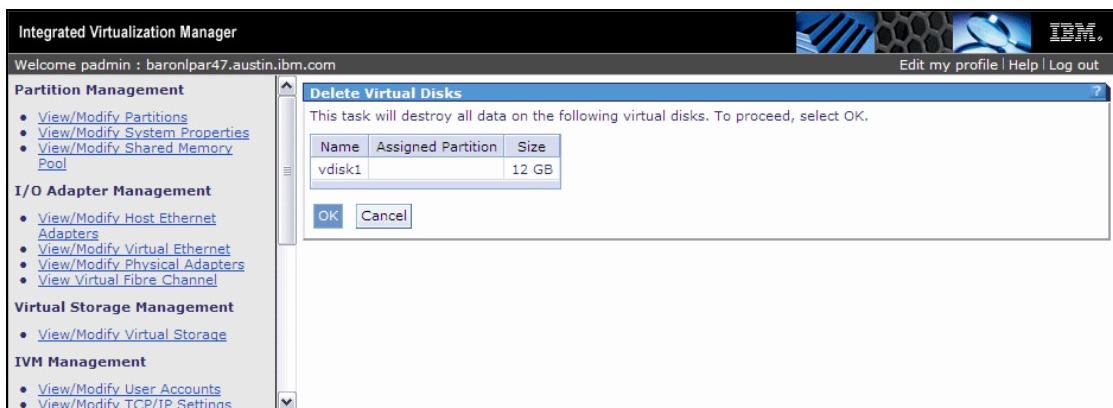


Figure 4-43 Confirm deletion of the virtual disk

The virtual disk will be deleted and the occupied space in the storage pool will become available.

4.6.4 Optical and Tape devices

Optical devices are CD or DVD drives. There are two types of optical devices:

- ▶ Physical optical devices
 - Local in BladeCenter media tray
 - Remote media on local desktop or laptop

- ▶ Virtual optical devices

Physical tape devices must be Serial Attached SCSI (SAS)

Physical optical devices

Physical optical devices are the CD or DVD drives installed in the media tray of a IBM BladeCenter. Each type of BladeCenter chassis is delivered with a CD drive or a DVD- drive.

The other physical optical device that can be used is remote media. An ISO image or an CD or DVD in your laptop or desktop can be assigned to the blade. The Web interface of the Advanced Management Module provides this capability.

As Table 4-1 shows, these two different drives can be identified by their location paths.

Table 4-1 Optical drive location paths

Location path	Description
U78A5.001.WIH01AA-P1-T1-L1-L2-L3	CD or DVD drive in the Media tray
U78A5.001.WIH01AA-P1-T1-L1-L1	Remote media

The name of the optical drives can vary, depending on the kind of drive or remote media you are using.

Before the BladeCenter physical optical device can be used, the media tray must be assigned to the blade slot you are working with. The physical optical device cannot be shared between LPARs or the VIOS and the LPARs.

The assignment of the physical optical device can be changed at any time. The assignment can be made or changed from the Optical Devices tab in the View/Modify Virtual Storage window. This section describes how to use the Storage Management to change the assignment of physical optical devices.

Changing the assignment of physical optical drives

To change the assignment of physical optical drives, click the **Optical/Tape** tab. Figure 4-44 on page 134 shows a list of available physical devices. The table in that figure indicates the physical optical device and shows the LPAR assignment. The example shows that cd1 is assigned to the partition named JS231par2.

The list of physical devices may vary, depending on the media tray assignment and the usage of remote media. Figure 4-44 on page 134 shows two physical CD-ROM drives. The second drive is a remote CD-ROM drive that is mounted

over the remote control interface of the Advanced Management Module (AMM) in the BladeCenter chassis.

Note: The remote control function for the IBM BladeCenter JS23 or JS43 is only available to the blade slot that has the media tray assignment.

To change the assignment of a physical optical device, select the check box of the device to be changed and click **Modify partition assignment**. A dialog opens that guides you through the assignment change.

The screenshot shows the IVM interface with the following details:

- Left Sidebar:** Lists various management categories with their sub-links:
 - Partition Management:** View/Modify Partitions, View/Modify System Properties
 - I/O Adapter Management:** View/Modify Host Ethernet Adapters, View/Modify Virtual Ethernet, View/Modify Physical Adapters, View Virtual Fibre Channel
 - Virtual Storage Management:** View/Modify Virtual Storage
 - IVM Management:** View/Modify User Accounts, View/Modify TCP/IP Settings, Guided Setup, Enter PowerVM Edition Key
 - System Plan Management:** Manage System Plans
 - Service Management:** Electronic Service Agent, Service Focal Point, Manage Serviceable Events, Service Utilities (Create Serviceable Event, Manage Dumps), Collect VPD Information, Updates, Backup/Restore, Application Logs, Monitor Tasks, Hardware Inventory
- Main Content Area:** Titled "View/Modify Virtual Storage".
 - Tab Bar:** Virtual Disks, Storage Pools, Physical Volumes, **Optical/Tape**
 - Physical Optical Devices:** A table showing physical optical devices.

Select	Name	Description	Assigned Partition	Physical Location Code
<input type="checkbox"/>	cd0	USB DVD-COMBO Drive	U78A5.001.W1H2302-P1-T1-L1-L2-L3	
<input type="checkbox"/>	cd1	USB DVD R/RW or RAM Drive	JS23lpar2 (2)	U78A5.001.W1H2302-P1-T1-L1-L1
 - Virtual Optical Media:** (No devices)
 - Physical Tape Devices:** (No devices)

Figure 4-44 Physical optical and tape devices in IVM

To make the physical optical device available to the VIOS itself, select **None**. Otherwise, select the logical partition that the physical optical device will be assigned and then click **OK**, as shown in Figure 4-45 on page 135.

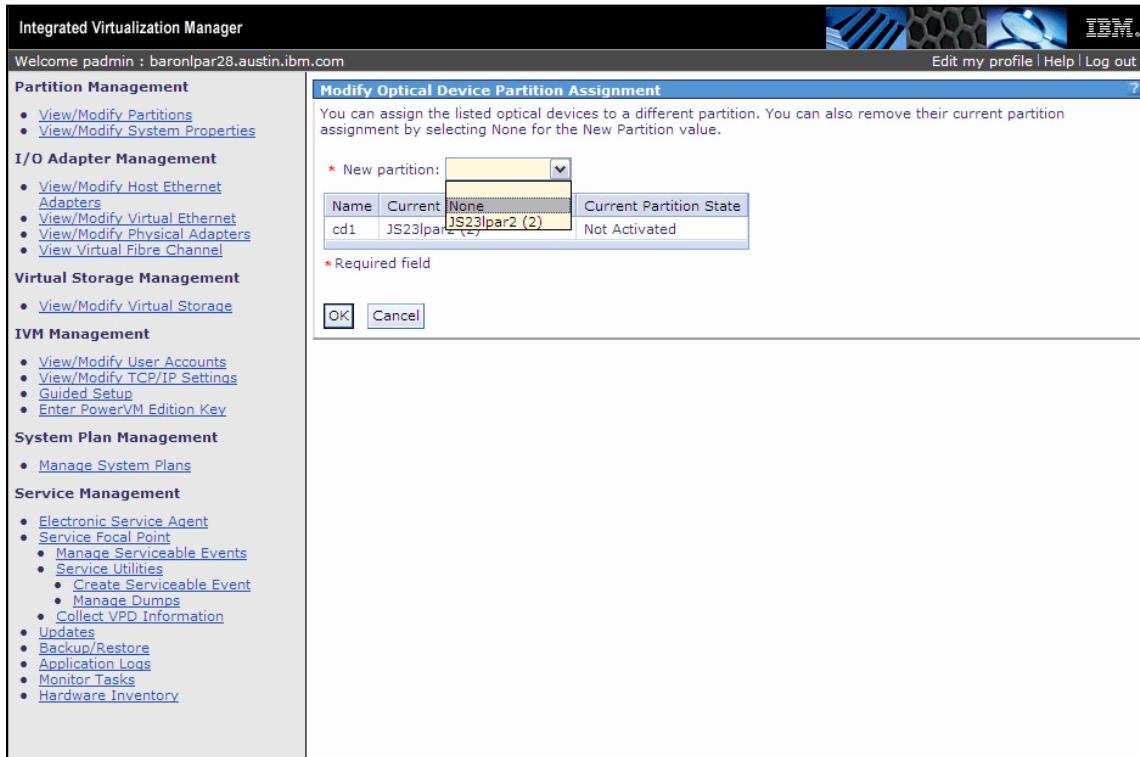


Figure 4-45 Change physical optical device assignment

Virtual optical devices

Virtual optical devices were introduced with Virtual I/O Server V1.5. Together with the Media Library of a Virtual I/O Server, this device is able to virtualize CD or DVD images that are stored in the VIOS media library to one or more logical partitions. Before virtual optical device can be used, you must configure a media library.

Creating a media library

To set up a media library, a storage pool must be available. Refer to 4.6.2, “Storage pools” on page 124, for an explanation about how to set up a storage pool.

To set up a media library, follow these steps:

1. Click the **Optical/Tape** tab in the View/Modify Virtual Storage window to create a media library. Then click the **Create Library** button, as shown in Figure 4-46 on page 136.

The screenshot shows the Integrated Virtualization Manager (IVM) interface. The left sidebar contains a navigation menu with several categories and their sub-links:

- Partition Management**: View/Modify Partitions, View/Modify System Properties, View/Modify Shared Memory Pool
- I/O Adapter Management**: View/Modify Host Ethernet Adapters, View/Modify Virtual Ethernet, View/Modify Physical Adapters, View Virtual Fibre Channel
- Virtual Storage Management**: View/Modify Virtual Storage
- IVM Management**: View/Modify User Accounts, View/Modify TCP/IP Settings, Guided Setup, Enter PowerVM Edition Key
- System Plan Management**: Manage System Plans
- Service Management**: Electronic Service Agent, Service Focal Point, Manage Serviceable Events, Service Utilities, Create Serviceable Event, Manage Dumps, Collect VPD Information, Updates, Backup/Restore, Application Logs, Monitor Tasks, Hardware Inventory

The main content area is titled "View/Modify Virtual Storage" and has tabs for Virtual Disks, Storage Pools, Physical Volumes, and Optical/Tape (which is selected). The Optical/Tape tab displays the following sections:

- Physical Optical Devices** (No devices)
- Virtual Optical Media** (No media library): A note states you can assign virtual optical media such as an ISO image directly to a partition for storage. It includes a "Create Library" button.
- Physical Tape Devices** (No devices)

Figure 4-46 Create media library

2. Select an available storage pool and the amount of storage space that will be used from this storage pool to create the media library, and then click **OK** as shown in Figure 4-47 on page 137.

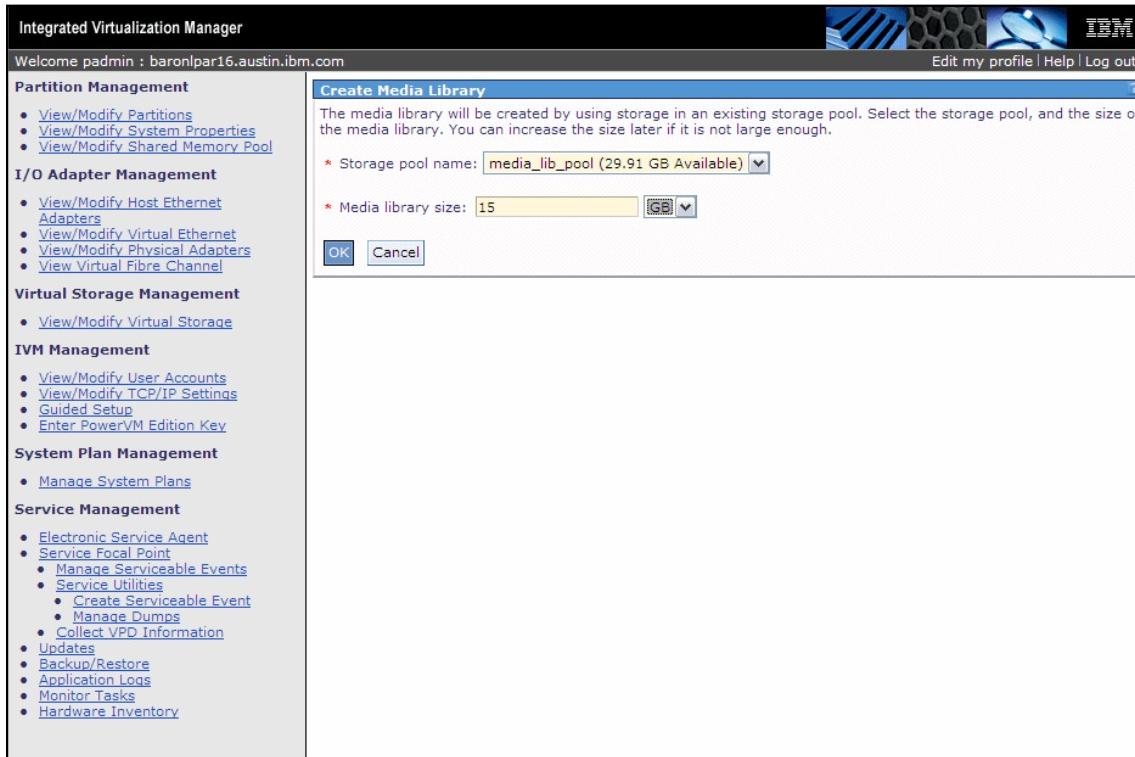


Figure 4-47 Media library size and storage pool

Depending on the size of the media library, the creation time will vary. After the media library is successfully created, the current view in the View/Modify Virtual Storage window will change, showing Media Library options.

The size of media library can be increased at any time by clicking the Extend Library button. Selecting the Delete Library button allows you to delete the complete media library and all added media and return the space to the storage pool.

Next, you need to add the new media into the media library.

Adding new media into the media library

New media can be added to the media library and later assigned to the virtual optical devices. The new media can consist of image files such as ISO images or copies from physical CDs or DVDs. In addition, you can create blank media that can be written to an LPAR using the virtual optical device.

To add new media in the media library, click **Add Media...** as shown in Figure 4-48 on page 138.

The screenshot shows the Integrated Virtualization Manager (IVM) web interface. The top navigation bar includes the title 'Integrated Virtualization Manager', the user 'padmin : baronpar16.austin.ibm.com', and links for 'Edit my profile | Help | Log out'. The main menu on the left has sections for 'Partition Management', 'I/O Adapter Management', 'Virtual Storage Management', 'IVM Management', 'System Plan Management', and 'Service Management'. The 'Virtual Storage Management' section is expanded, showing options for 'View/Modify Virtual Storage'. The right panel is titled 'View/Modify Virtual Storage' and has tabs for 'Virtual Disks', 'Storage Pools', 'Physical Volumes', and 'Optical/Tape' (which is selected). Under 'Optical/Tape', there are sections for 'Physical Optical Devices (No devices)' and 'Virtual Optical Media'. A note states: 'You can assign virtual optical media, such as an ISO image, directly to a partition to use for storage. Select the virtual optical media, then select the task that you want to perform. You can also extend the size of the media library or delete an existing media library.' Below this is a table with columns 'Select', 'Name', 'Assigned Partition', 'Mount Type', and 'Size'. At the bottom of the table is a 'Add Media...' button. The overall interface is a standard web-based administrative tool with a dark theme and blue highlights.

Figure 4-48 Add media to media library

There are four options to create new media:

- ▶ Upload media
- ▶ Add existing file
- ▶ Import from physical optical device
- ▶ Create blank media

The Upload media option allows you to transfer files or ISO images from a workstation directly to the media library. There is a limitation in the file size of 2 GB for this option.

The Add existing file option adds an existing file that is available in a VIOS file system as new media.

The Import from physical optical device option allows you to use a the physical CD or DVD. The data will be copied from the CD or DVD into the media library.

Note: Our testing revealed that the local CD or DVD drive in the media tray of the BladeCenter chassis is a faster option compared to the remote media option with a physical CD or DVD drive.

The Create blank media option allows you to create blank media that may be written to from an LPAR.

Figure 4-49 shows an example that uses Import from physical optical device to create the new media. Click **OK** to start the copy task.

Note: Do not use spaces in the name of the new media. If you use spaces in the name, IVM will return an error message stating there are too many parameters.

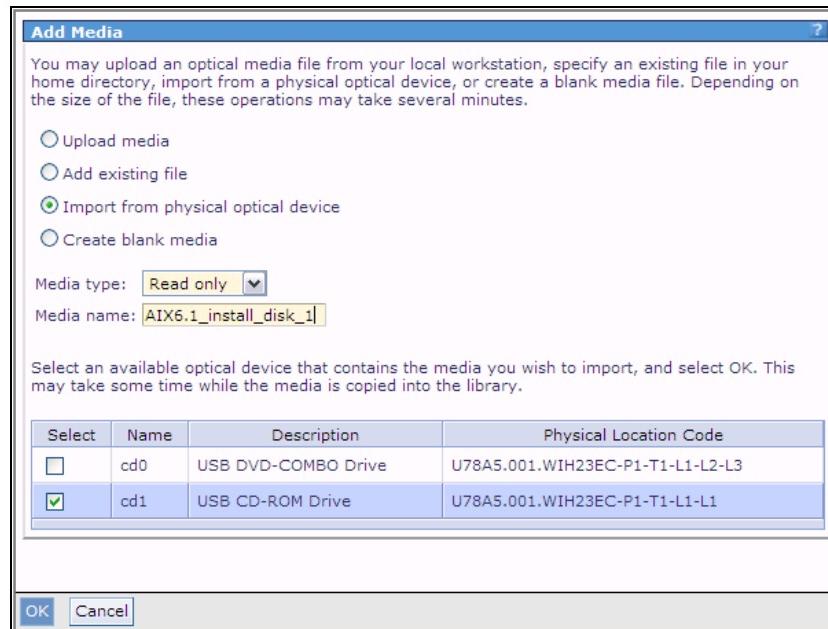


Figure 4-49 Add new media - settings

The copy task takes some time to complete. While the copy task is running, you may proceed with other configuration tasks after you see the message Performing Task - Please Wait displayed, as shown in Figure 4-50.

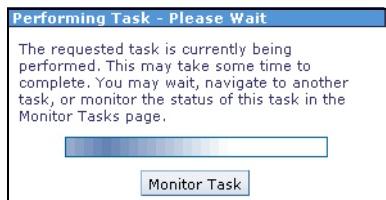


Figure 4-50 Performing task

Click the **Monitor Task** link from the Navigation area to verify the completion of the task. Monitor Tasks contains a list of events and the status, either running, successful, or failed.

Note: An alternative way to monitor the process of creating new media is to review the list under the Optical Devices tab, as shown in Figure 4-51 on page 141.

If your new media is not listed here, click the **Refresh** button. During the copy operation, the new media is shown as Read/Write and the size will increase on a refresh. After the copy operation is finished, the mount type will change to **Read only**.

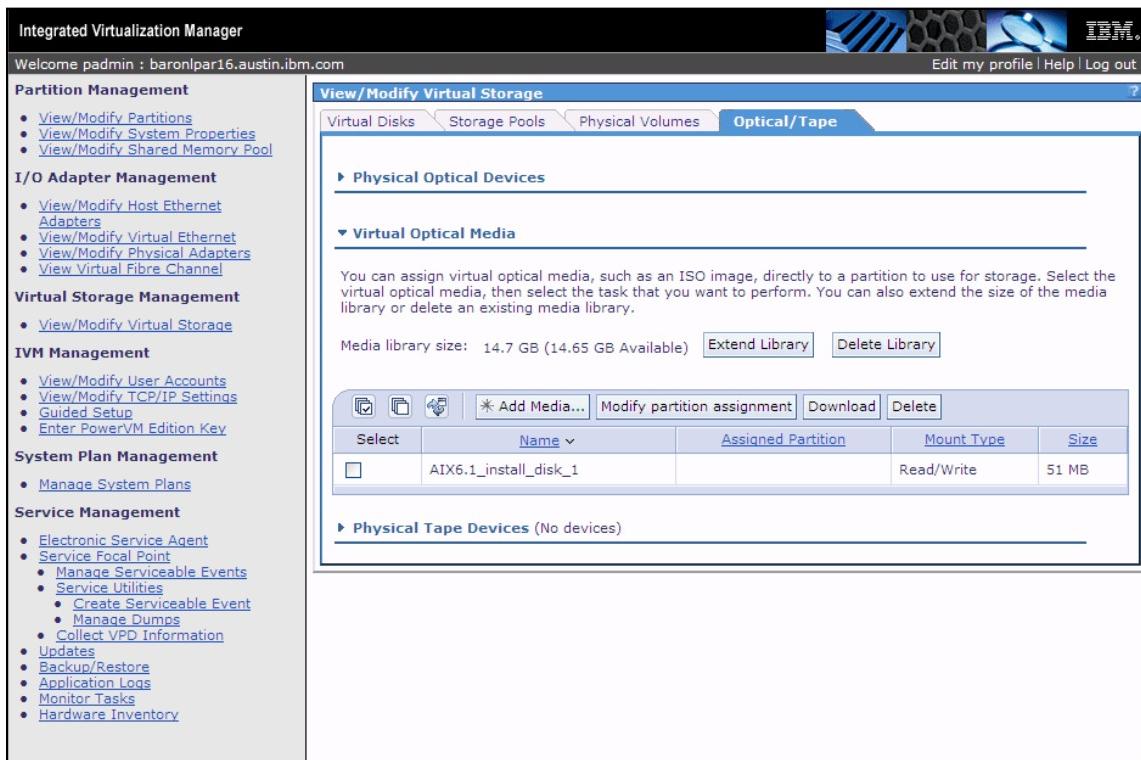


Figure 4-51 Newly created media with the copy operation in progress

Modifying media assignment to virtual optical devices in logical partitions

Media can be assigned from the Optical/Tape tab in the View/Modify Virtual Storage window, when using the Create Partition wizard or from the Partition Properties window. The next step will be to modify the partition assignment of the media in the media library.

Note: The logical partition *must* have a virtual optical device assigned prior to assigning or modifying the partition assignment of the selected media.

To modify the virtual media assignment, select the desired media name then click **Modify partition assignment** under the Optical/Tape tab from the View/Modify Virtual Storage window.

As shown in Figure 4-52 on page 142, the media AIX6.1_install_disk_1 is not assigned to any LPAR. This media since it is read only will be selected and

assigned to two LPARs. Select the check box for the desired media and click the **Modify partition assignment** button.

The screenshot shows the Integrated Virtualization Manager web interface. The left sidebar contains navigation links for various management categories like Partition Management, I/O Adapter Management, and Service Management. The main content area is titled 'View/Modify Virtual Storage' and has tabs for Virtual Disks, Storage Pools, Physical Volumes, and Optical/Tape (which is selected). Under 'Virtual Optical Media', it says: 'You can assign virtual optical media, such as an ISO image, directly to a partition to use for storage. Select the virtual optical media, then select the task that you want to perform. You can also extend the size of the media library or delete an existing media library.' Below this is a table:

Select	Name	Assigned Partition	Mount Type	Size
<input checked="" type="checkbox"/>	AIX6.1_install_disk_1		Read only	621 MB

Figure 4-52 Modify partition assignment

As shown in Figure 4-53 on page 143 no LPARs are assigned to the media AIX6.1_install_disk_1. Next, LPARs JS23DM1par4 and JS23 DP1par5 will be assigned the same media by selecting the check box next to the logical partitions. Choose the Media type **Read only** or **Read/Write** and click **OK**. Only Read only media can be assigned to more than one LPAR.

The screenshot shows the IVM interface with a sidebar containing various management links. The main area is titled 'Modify Media Partition Assignment' with a help icon and a question mark icon.

Modify Media Partition Assignment

You can modify the partitions to which the media is assigned by selecting the appropriate virtual optical devices. Read-only media may be assigned to more than one device.

Only partitions containing virtual optical devices are listed. Use the optical tab in the partition properties task to create virtual optical devices.

Media name: AIX6.1_install_disk_1

Media type:

Select	Partition	Device	Current Media	Current Partition State
<input checked="" type="checkbox"/>	JS23DM1par4 (4)	vtopt0	None	Running
<input type="checkbox"/>	JS23AMSlpar3 (3)	vtopt1	None	Running
<input checked="" type="checkbox"/>	JS23DP1par5 (5)	vtopt2	None	Not Activated

Figure 4-53 Modify media partition assignment

Click **OK** to return to the view of the optical devices. Notice that the updated table shown in Figure 4-54 on page 144 now contains the LPARs JS23DM1par4 and JS23_DP1par5 in the Assigned Partition column as assigned partitions for the media AIX6.1_install_disk_1.

Figure 4-54 New assigned media to partitions

A media can be removed from a partition following the same procedure by deselecting the media that is assigned to the partition.

4.7 Partition configuration for Virtual I/O Client (VIOC)

With networking and storage defined, you can now create additional VIOC LPARs for the installation of additional supported operating systems.

4.7.1 Live Partition Mobility considerations

If Live Partition Mobility (LPM) is going to be considered for a VIOC in logical partitions, you need to review the following planning considerations (currently unavailable on IBM i).

- ▶ VIOS running on source and target IBM BladeCenter JS23 or JS43:
 - VIOS should be at the latest fixpack.

- IBM BladeCenter JS23 or JS43 should be at the latest system firmware.
- ▶ All I/O must be virtual to the LPAR:
 - SEA adapters are required. No HEA logical ports can be assigned.
 - No virtual optical drives can be assigned.
 - No physical adapters can be assigned.
- ▶ SAN storage properly configured for sharing between the two Virtual I/O Servers.
- ▶ Processor compatibility modes between source and target systems.
- ▶ Memory region sizes must match between source and target systems.
- ▶ If Active Memory Sharing (AMS) is being used on the source VIOS/VIOC, it must be available on the target VIOS.
- ▶ Only IVM- to-IVM managed systems are allowed (no HMC-to-IVM or IVM-to-HMC is allowed)

Chapter 11, “Performing Live Partition Mobility” on page 433 provides setup and configuration details for Partition Mobility with a JS23 or JS43. The Redbooks publication *PowerVM Live Partition Mobility on IBM System p*, SG24-7460, provides an overview of Partition Mobility and explains how it applies to other System p configurations.

4.7.2 Partition name and environment

The create partition wizard is started with the **Create Partition** button on the **View/Modify Partitions** link in the navigation area of the IVM UI, as shown in Figure 4-55 on page 146.

Integrated Virtualization Manager

Welcome padmin : baronpar47.austin.ibm.com

Partition Management

- [View/Modify Partitions](#)
- [View/Modify System Properties](#)
- [View/Modify Shared Memory Pool](#)

I/O Adapter Management

- [View/Modify Host Ethernet Adapters](#)
- [View/Modify Virtual Ethernet](#)
- [View/Modify Physical Adapters](#)
- [View Virtual Fibre Channel](#)

Virtual Storage Management

- [View/Modify Virtual Storage](#)

IVM Management

- [View/Modify User Accounts](#)
- [View/Modify TCP/IP Settings](#)
- [Guided Setup](#)
- [Enter PowerVM Edition Key](#)

System Plan Management

- [Manage System Plans](#)

Service Management

- [Electronic Service Agent](#)
- [Service Focal Point](#)
 - [Manage Serviceable Events](#)
 - [Service Utilities](#)
 - [Create Serviceable Event](#)
 - [Manage Dumps](#)
 - [Collect VPD Information](#)
- [Updates](#)
- [Backup/Restore](#)
- [Application Logs](#)
- [Monitor Tasks](#)
- [Hardware Inventory](#)

View/Modify Partitions

To perform an action on a partition, first select the partition or partitions, and then select the task.

System Overview

Total system memory:	8 GB	Total processing units:	8
Memory available:	2.22 GB	Processing units available:	5.3
Reserved firmware memory:	704 MB	Processor pool utilization:	0.16 (2.0%)
Available shared memory pool size:	768 MB		
System attention LED:	Inactive		

Partition Details

Select	ID	Name	State	Uptime	Memory	Processors	Entitled Processing Units	Utilized Processing Units	Reference Code
<input type="checkbox"/>	1	is43-vios	Running	4.78 Days	1.5 GB	8	0.8	0.12	
<input type="checkbox"/>	2	rhel53	Running	3.14 Days	1 GB	1	0.1	0.03	Linux ppc64
<input type="checkbox"/>	3	IBMi	Not Activated		1 GB	1	1.0		00000000
<input type="checkbox"/>	4	mobilepar	Running	4.75 Days	608 MB	8	0.8	0.02	

Figure 4-55 View/Modify Partition

4.7.3 Partition name

When the wizard starts, a new window will open as shown in Figure 4-56 on page 147. This gives you the opportunity to change the Partition ID number, provide a Partition name, and select an operating system environment. Select the **Next** button for the memory step.

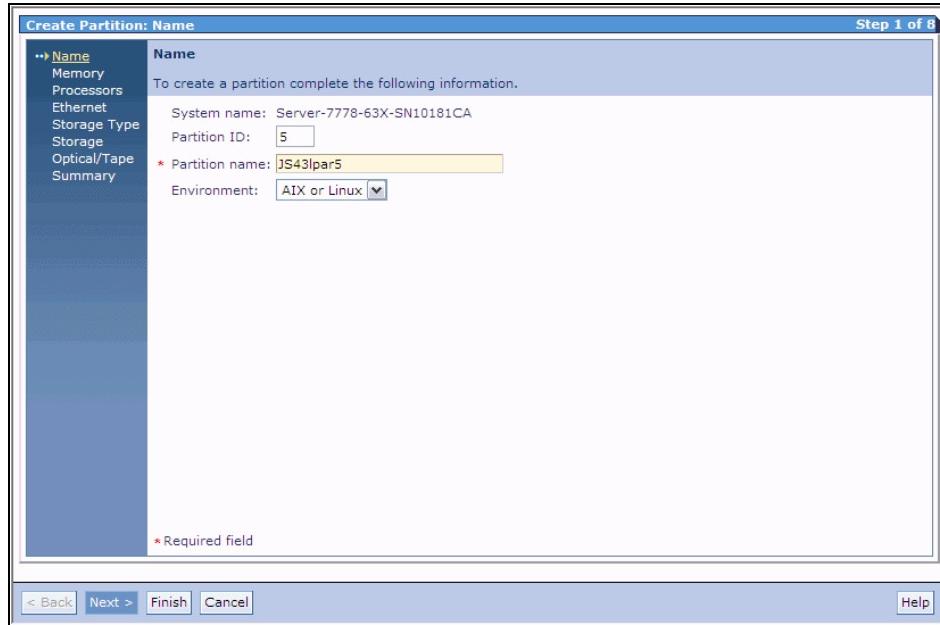


Figure 4-56 Create Partition: Name

4.7.4 Partition Memory

Figure 4-57 on page 148 shows how to assign memory to the partition. The two memory options are dedicated and shared. In this section we will only discuss dedicated memory. Shared memory is covered in Chapter 5, “Active Memory Sharing configuration using IVM” on page 177.

Total system memory and the current memory available for a new partition is summarized under memory mode selection section. The amount entered in the box will determine the initial *Assigned* and *Maximum* values in the LPAR partition properties. After you enter the desired amount of memory, select the **Next** button.

Note: IVM does not allow you to over-commit *dedicated* memory resources.

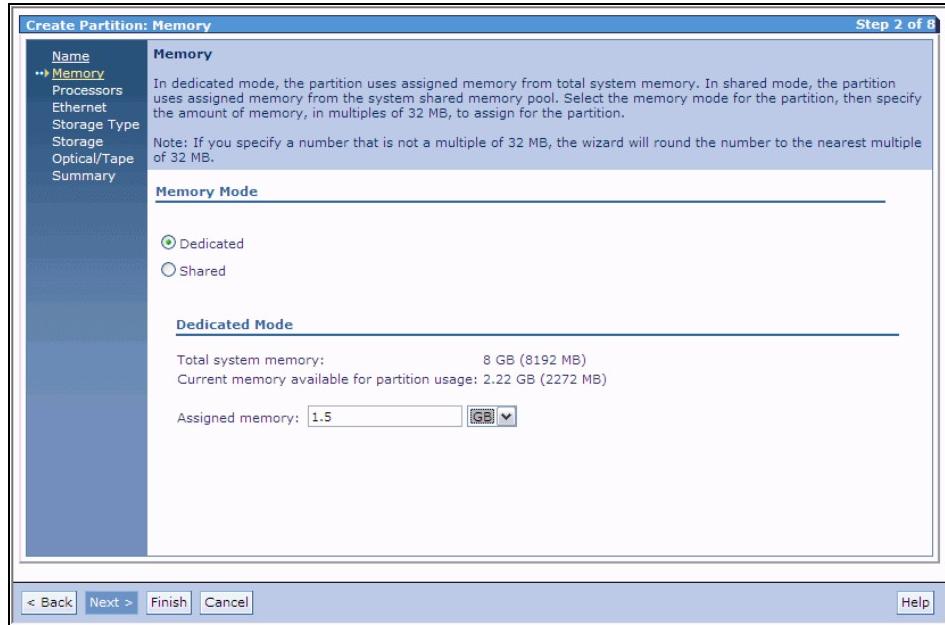


Figure 4-57 Create Partitions: Memory

4.7.5 Partition Processors

On the Create Partition: Processors window you have the option of assigning dedicated or shared processors. In shared mode, for each virtual processor, 0.1 processing units will be assigned. In dedicated mode, each assigned processor uses one physical processor.

Available processor resources are displayed on the window and, as with dedicated memory resources, they cannot be over-committed. Figure 4-58 on page 149 shows a selection of shared mode and eight assigned processors for this example. After you make your selections, click the **Next** button.

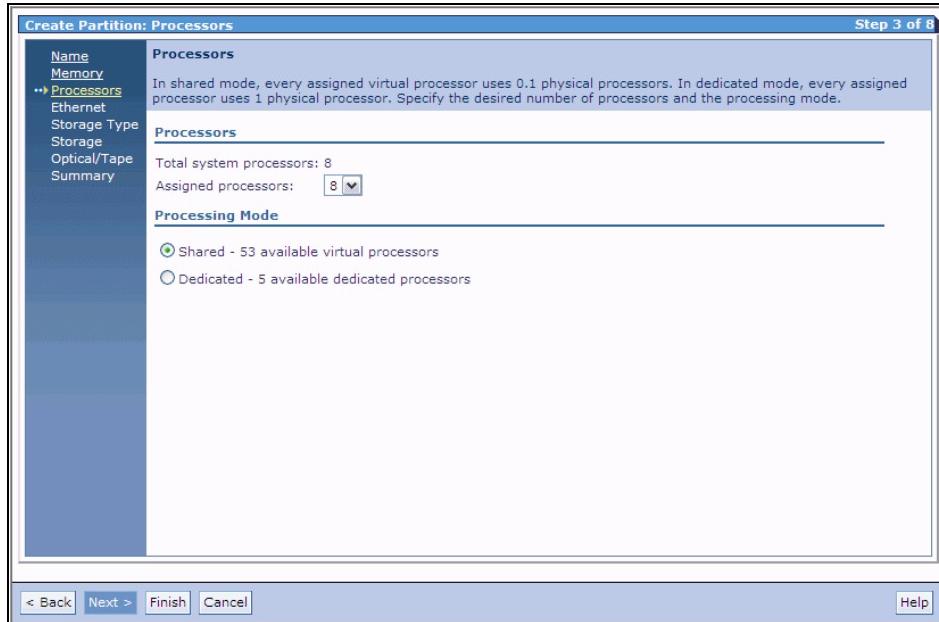


Figure 4-58 Create Partition: Processors

Note: After an LPAR is created, the processor mode cannot be changed from shared to dedicated or dedicated to shared from IVM, only from the VIOS CLI using the `chsyscfg` command.

4.7.6 Partition Ethernet

The Create Partition: Ethernet window displays the choices for assigning network connectivity. The choices, as previously noted, are:

- ▶ HEA logical port
- ▶ Virtual adapter
- ▶ SEA adapter
- ▶ Physical adapter

HEA logical ports, Virtual adapters, or SEA adapters can be selected during this step. Physical adapters, if available, are assigned during a later step in the Create partition wizard.

Note: If creating a LPAR with share memory resources you only have the option of assigning virtual Ethernet adapters.

Figure 4-59 shows the first three options. The selection in this example is virtual Ethernet adapter 1 on the logical partition assigned to a SEA adapter. Note that you also have an opportunity as this time to create additional virtual Ethernet adapters for the logical partition.

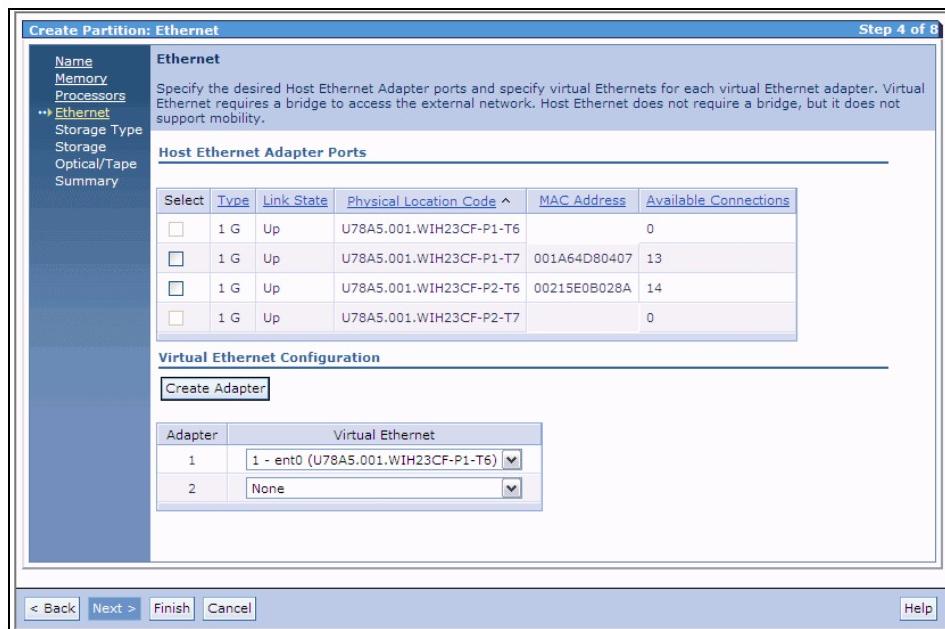


Figure 4-59 Create Partition: Ethernet

Note: HEA logical ports and physical adapter assignments cannot be used on logical partitions that will be considered for Partition Mobility.

4.7.7 Partition Storage Types and Assignments

Logical partition disk storage can be virtual disks from a storage pool or physical volumes. Figure 4-60 on page 151 shows these choices, and also provides the option to create additional virtual disks from a storage pool that has already been defined. Refer to 4.6.2, “Storage pools” on page 124 for an explanation about how to create storage pools.

Note: For logical partitions that will be used in Partition Mobility LPARs, the storage type must be physical volumes.

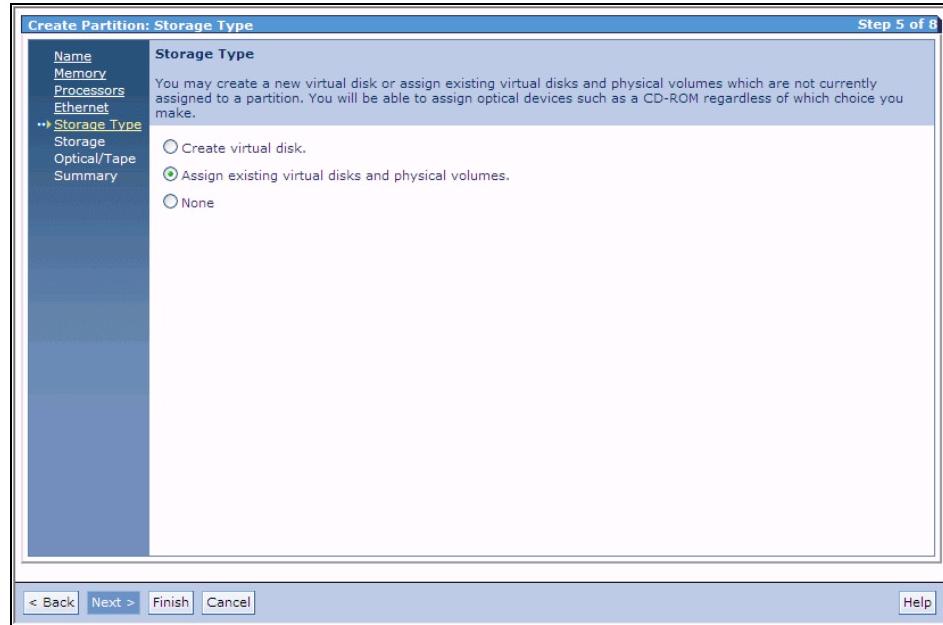


Figure 4-60 Create Partition: Storage Type

In this example we are using physical volumes. Click the option **Assign existing virtual disks and physical volumes**, and then click **Next**.

Figure 4-61 on page 152 shows the available physical volumes. Note that no virtual disks have been defined for this example, so the table under Available Virtual Disks is empty. Select one or more available hdisks, then click the **Next** button.

Note: For initial LPAR setup and OS install, we recommend that you only select the install target disk at this time.

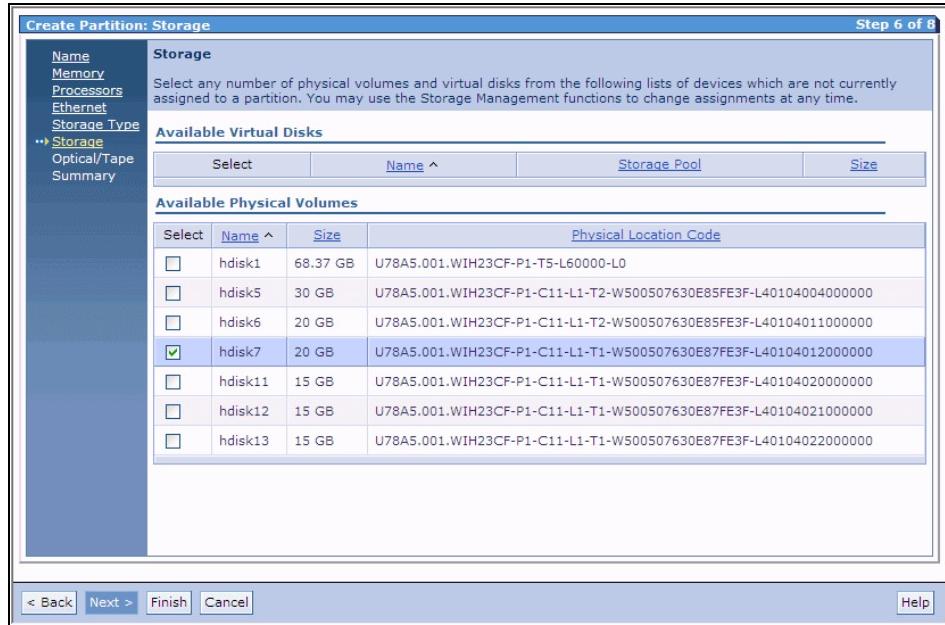


Figure 4-61 Logical Partition: Storage

4.7.8 Optical and tape devices

Optical devices, both physical and virtual and physical tape devices, can be assigned to an LPAR. With an IBM BladeCenter JS23 or JS43, the physical optical device must be available to the BladeCenter slot that you are working with through the media tray assignment before assignment to an LPAR can be made.

Virtual Optical Devices are not dependent on the media tray assignment. Refer to “Virtual optical devices” on page 135 for an explanation of how to create the media library and virtual optical devices.

Note: Physical and virtual optical devices cannot be used on logical partitions that will be used in Partition Mobility.

A virtual tape or virtual optical device is required for backup for IBM i.

As of this writing, only the IBM System Storage TS2240 SAS attached tape system is supported in a BladeCenter JS23/43 VIOS environment.

Figure 4-62 on page 153 shows the optical device selection window. In this example, no physical optical device are available. By default the LPAR wizard will

present and select a virtual optical device. If a virtual optical device is not desired uncheck the selection box.

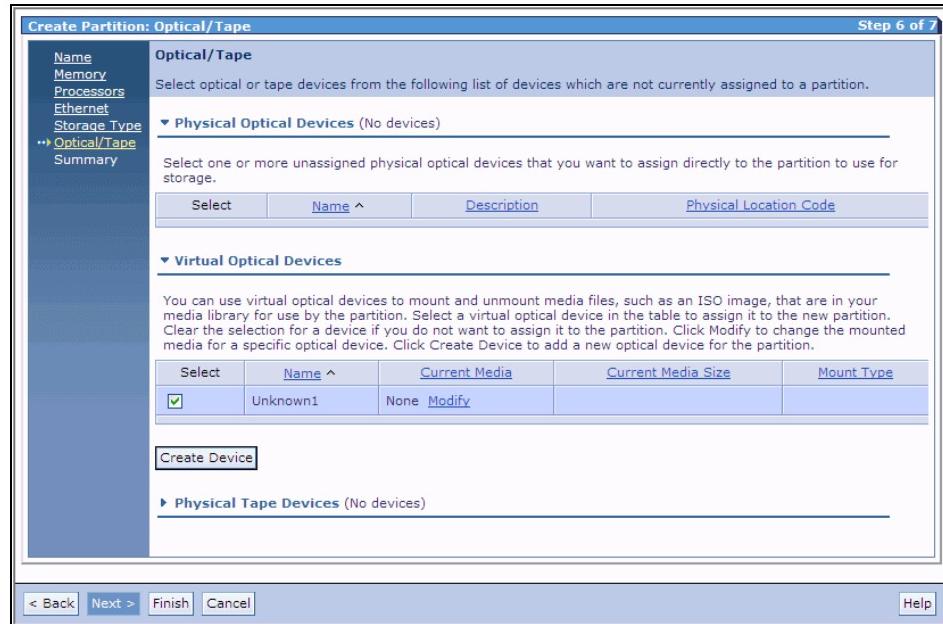


Figure 4-62 Create Partition: Optical

If unassigned physical adapters are available on the system the next window will provide the opportunity to assign them to the LPAR being created. If no physical adapter resources are available you will be directed to the summary window. Click the **Next** button to proceed to the Physical Adapters window (if available) or the Summary window.

4.7.9 Physical adapters

The physical adapter step will only be shown if I/O adapters are available for assignment to an LPAR. Figure 4-63 on page 154 shows availability of an Ethernet adapter that could be assigned to the example LPAR.

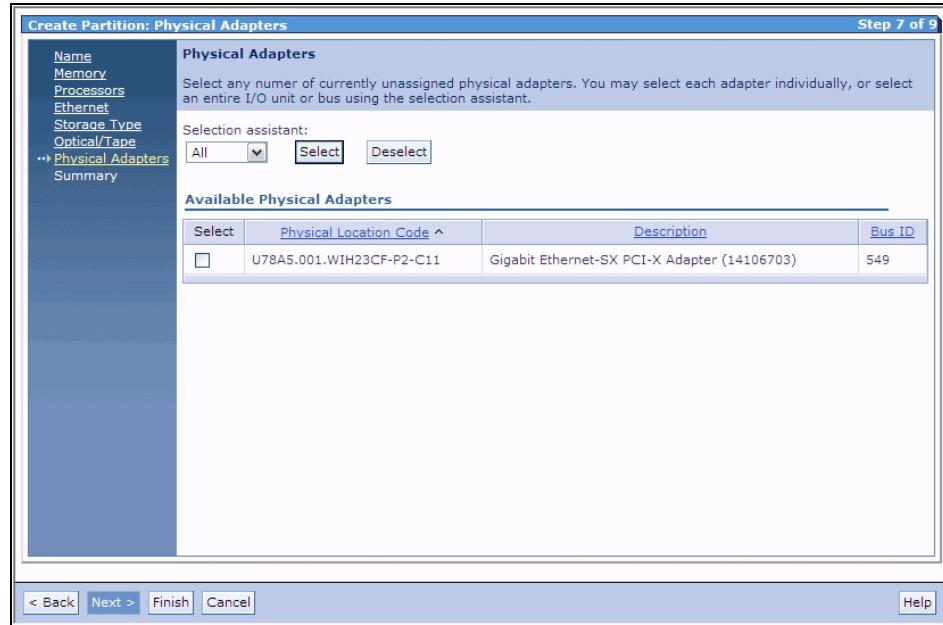


Figure 4-63 Create Partition: Physical Adapters

4.7.10 Partition Summary

The final window of the Create Partition wizard is the Create Partition: Summary, as shown in Figure 4-64 on page 155. All of the previous selections can be reviewed on this window and edited if required by using the **Back** button.

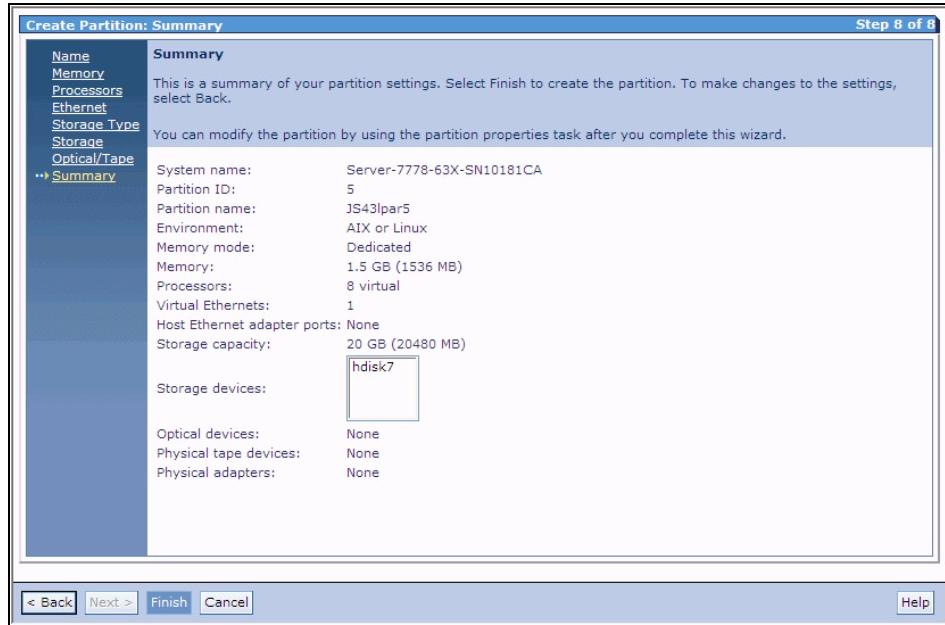


Figure 4-64 Create Partition: Summary

After your review is done and any needed adjustments have been made, click the **Finish** button to complete the logical partition creation.

Figure 4-65 on page 156 of the View/Modify Partitions window shows the new logical partition that was created.

Select	ID	Name	State	Uptime	Memory	Processors	Entitled Processing Units	Utilized Processing Units	Reference Code
<input type="checkbox"/>	1	js43-vios	Running	4.78 Days	1.5 GB	8	0.8	0.06	
<input type="checkbox"/>	2	rhel53	Running	3.14 Days	1 GB	1	0.1	0.03	Linux ppc64
<input type="checkbox"/>	3	IBMi	Not Activated		1 GB	1	1.0		00000000
<input type="checkbox"/>	4	mobilelpar	Running	4.75 Days	608 MB	8	0.8	0.02	
<input type="checkbox"/>	5	JS43lpar5	Not Activated		1.5 GB	8	0.8		00000000

Figure 4-65 View/Modify Partition showing new partition

4.7.11 Partition properties changes and DLPAR operations

The IVM UI provides quick access to change an LPAR's properties and perform Dynamic LPAR (DLPAR) operations on an active LPAR. The IBM BladeCenter JS23 or JS43 have the capability to perform DLPAR operations on memory, processors, and real or virtual I/O adapters.

Partition properties changes and DLPAR operations are accessed by clicking the name of the partition to be modified when viewed from View/Modify Partitions view. When a partition name is selected, a new window will open to the General tab as shown in Figure 4-66 on page 157.

You can change the partition name, attention LED state, boot mode, and keylock position. You can also set participation in a partition workload group, as well as how the LPAR will start in the event of a complete managed system restart. The default setting is for the LPAR to automatically start when the system starts.

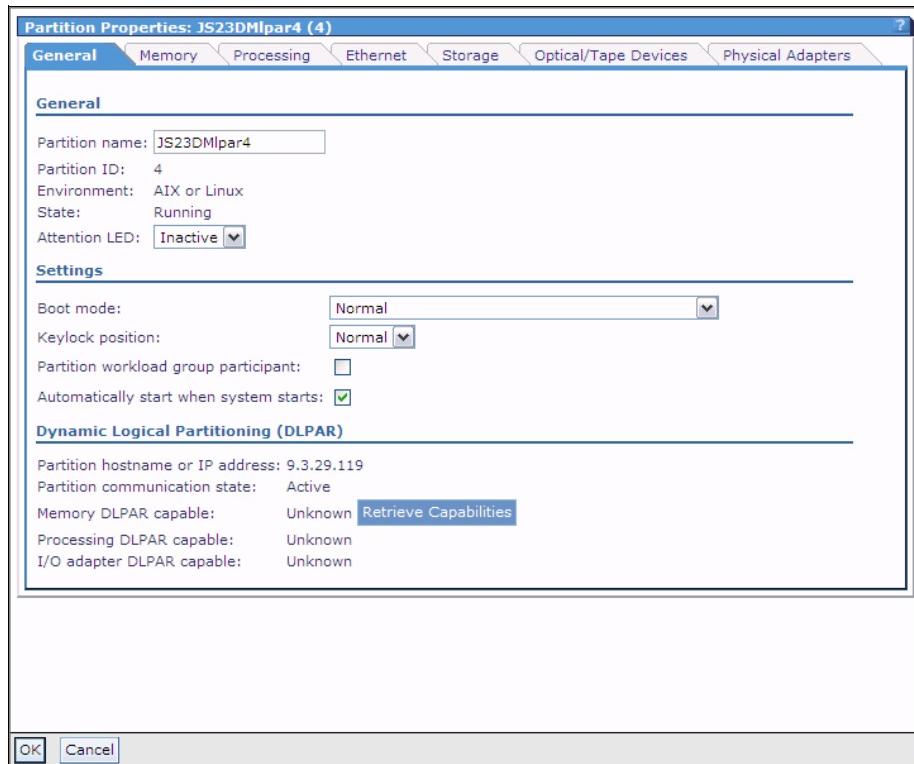


Figure 4-66 Partition Properties General tab

DLPAR capabilities can be retrieved by clicking the **Retrieve Capabilities** button. Figure 4-67 on page 158 shows the DLPAR capabilities of the IBM BladeCenter JS23 or JS43.

IBM i LPARs have a different Partition Properties General tab view. See 7.3, “Creating an IBM i 6.1 partition” on page 271 for more information.

Note: The LPAR must have a network interface configured and have an active Resource Monitoring and Control (RMC) connection with the management partition you can retrieve capabilities or perform DLPAR operations.

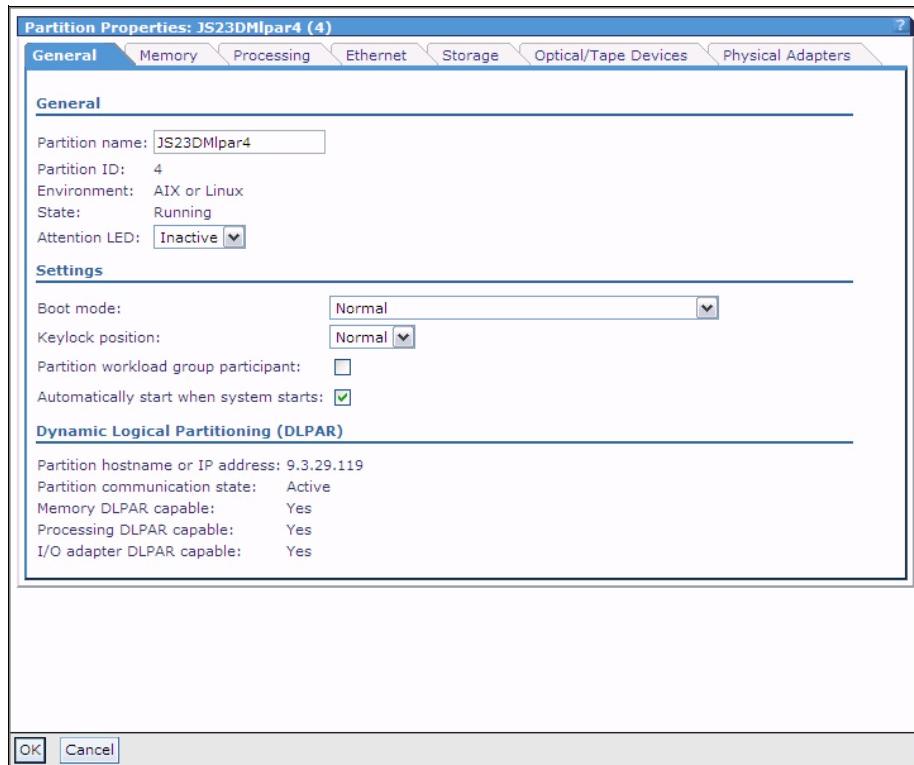


Figure 4-67 DLPAR retrieved capabilities

Selecting the **Memory** tab will display current and pending memory values for the LPAR, as shown in Figure 4-68 on page 159. In addition if a shared memory pool has been configured you will have the option to change between dedicated and shared memory. The change between dedicated and shared can only be done on an inactive LPAR.

An active LPAR can have its *Assigned* memory value changed between the range of the minimum and maximum values as a DLPAR operation. The Minimum memory and Maximum memory pending values can only be changed when the LPAR is *not* activated.

Note: The VIOS or management partition can change the Minimum and Maximum memory values while active, but will not become effective until a partition shutdown and restart is performed.

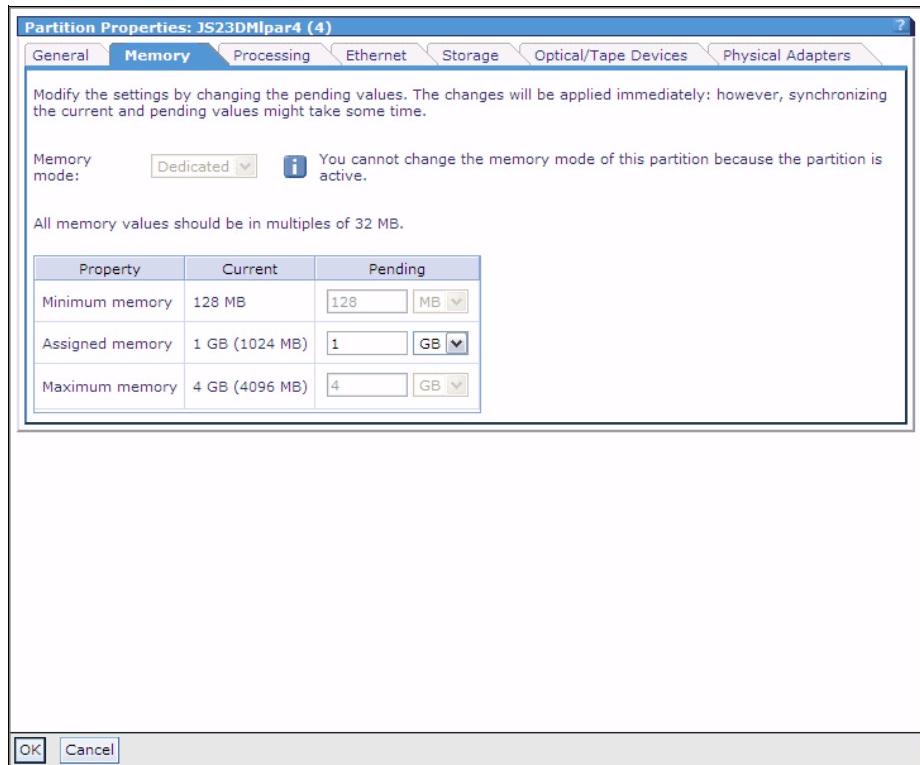


Figure 4-68 Partition Properties Memory tab

The **Processing** tab is used to change the processing units, virtual processors, partition priority weighting, and processor compatibility mode for LPARs using a shared processor pool, as shown in Figure 4-69. When changing the processor compatibility mode, a partition shutdown and restart is required for an active LPAR to make the change. If the LPAR is already inactive an activation is required before the current value will be updated.

Note: For Partition Mobility processor compatibility mode must be checked when moving an LPAR from a JS23 or JS43 to a JS12 or JS22. The earlier blades do not support Power6+ or Power6+ Enhanced modes.

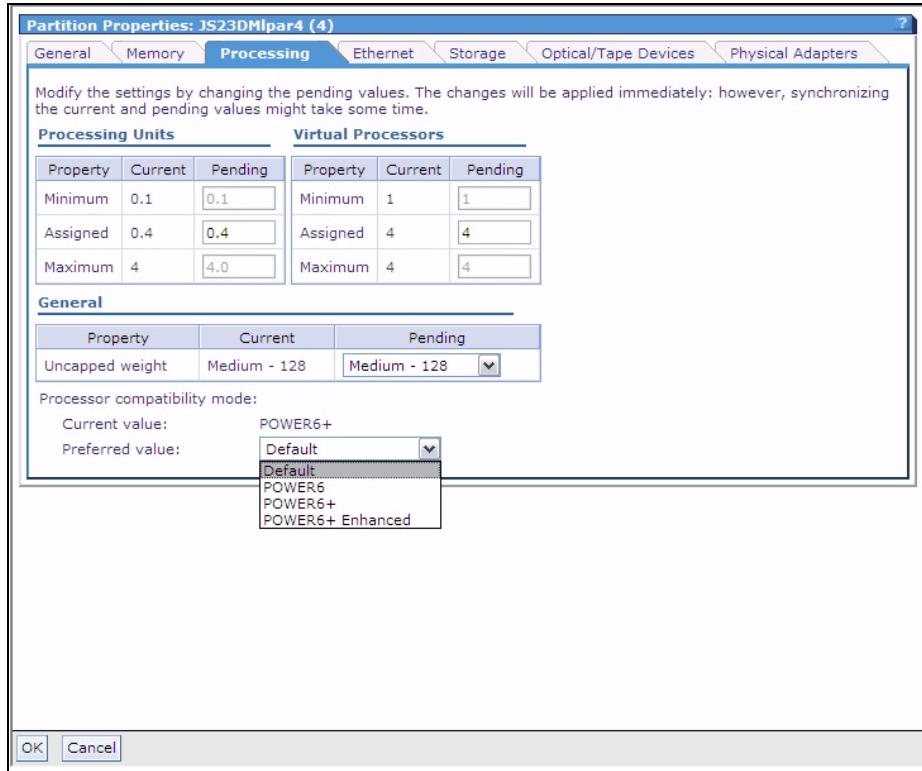


Figure 4-69 Partition Properties, Processing tab for shared pool

Partitions using dedicated processors will display the window as shown in Figure 4-70 on page 161. This example shows the LPAR in a not activated state and the minimum, assigned, and maximum values can be changed. In an active LPAR, only the assigned value can be altered as a DLPAR operation.

This window also allows changing the mode of sharing idle processors. The four modes are

- ▶ When inactive (default), shares excess capacity when the LPAR is not active
- ▶ When active, shares excess capacity when the LPAR is active, but not when inactive
- ▶ Always, shares excess capacity when the LPAR is active or inactive
- ▶ Never

The excess capacity that is shared is utilized by the shared pool.

The processor compatibility mode can also be changed when using dedicated processors.

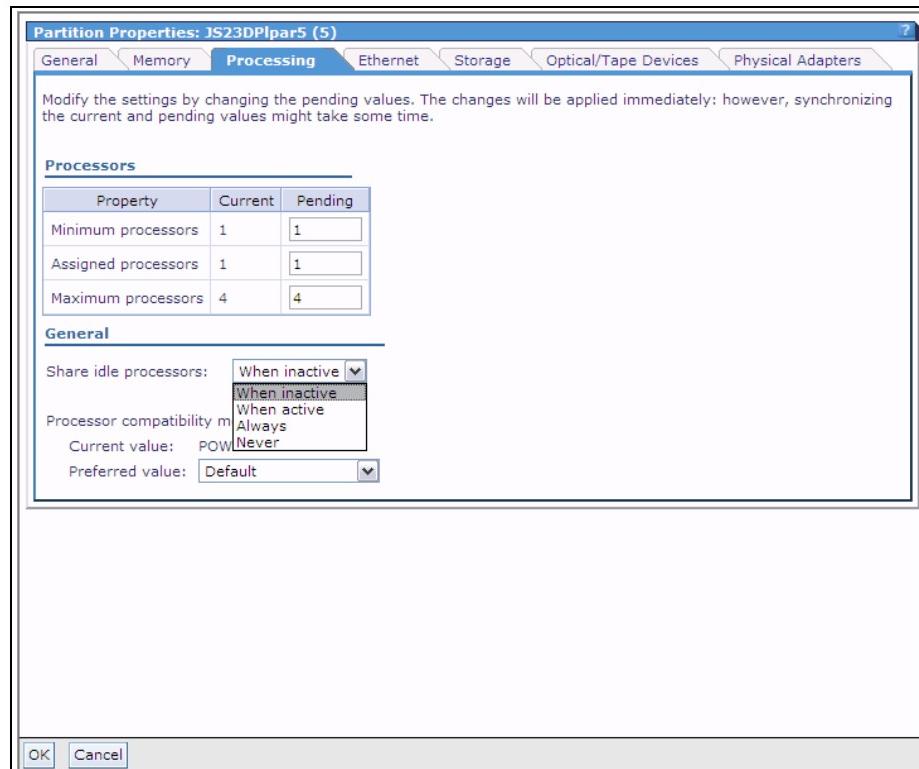


Figure 4-70 Partition Properties, Processing tab for dedicated processors

The **Ethernet** tab in Partition Properties allows the addition or removal of Ethernet adapters, as shown in Figure 4-71 on page 162.

Note: Before you can DLPAR remove Ethernet adapters from an active AIX LPAR, first use the **rmdev** command to removed the devices from the LPAR.

HEA virtual ports required the removal of Logical Host Ethernet Adapter (l-hea) and the Logical Host Ethernet Port (lp-hea). Virtual Ethernet adapters can be removed by deleting the Virtual I/O Ethernet Adapter (l-lan). Physical Ethernet adapters require the deletion of the adapter (ent) and the parent. The parent can be determined by the **lsdev** command. For example:

```
# lsdev -Cl ent1 -F parent  
pci1
```

The alternate method is to try the DLPAR command; in that case IVM will return an error message with details about the parent and child devices that must be removed.

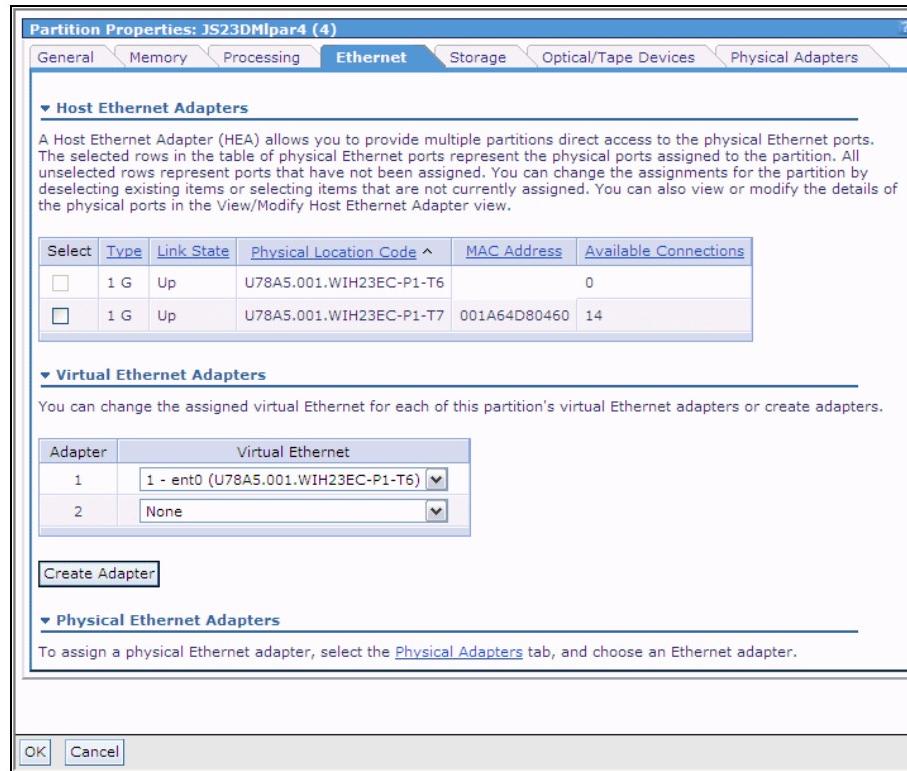


Figure 4-71 Partition Properties, Ethernet tab

Note: Partitions that are configured for shared memory or IBM i partitions cannot own HEAs. Therefore, the Host Ethernet Adapter section of this window will not be shown when viewing the properties of these types of LPARs.

The **Storage** tab can be used to add or remove storage devices, either physical volumes or virtual disks, as shown in Figure 4-72.

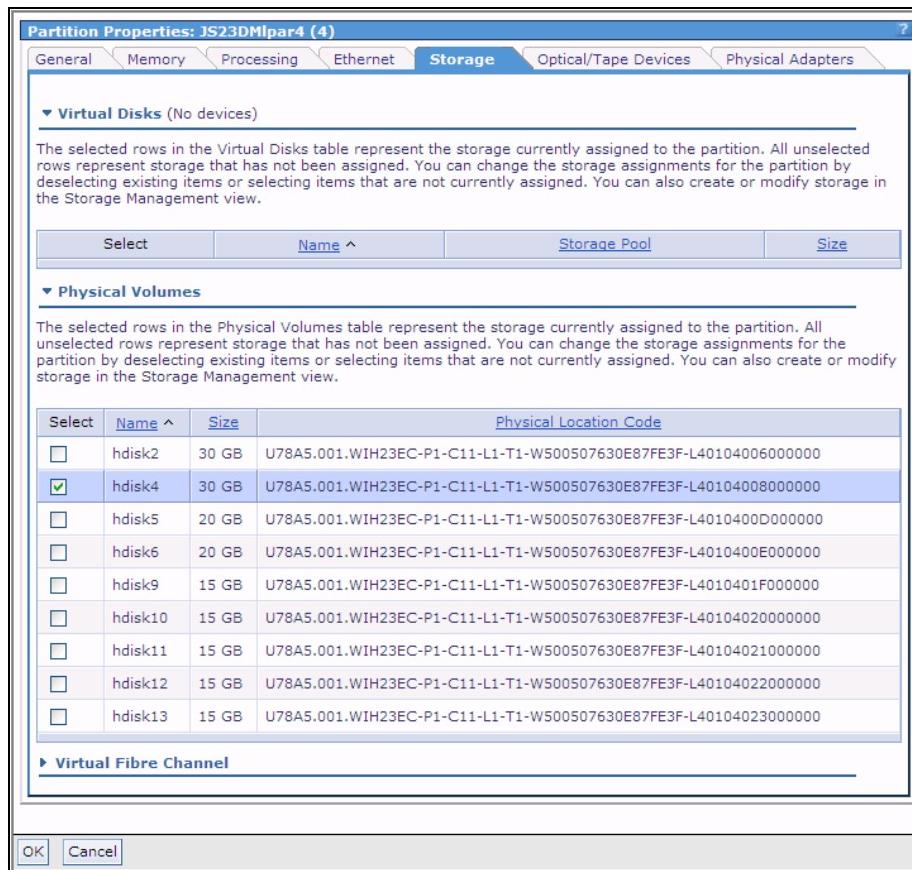


Figure 4-72 Partition Properties, Storage tab

Note: When removing storage from an active partition, IVM will require that you verify that a forced device removal from a running partition is requested.

Optical device assignments, both physical and virtual, and physical tape assignments can be managed from the **Optical /Tape Devices** tab shown in Figure 4-73 on page 164.

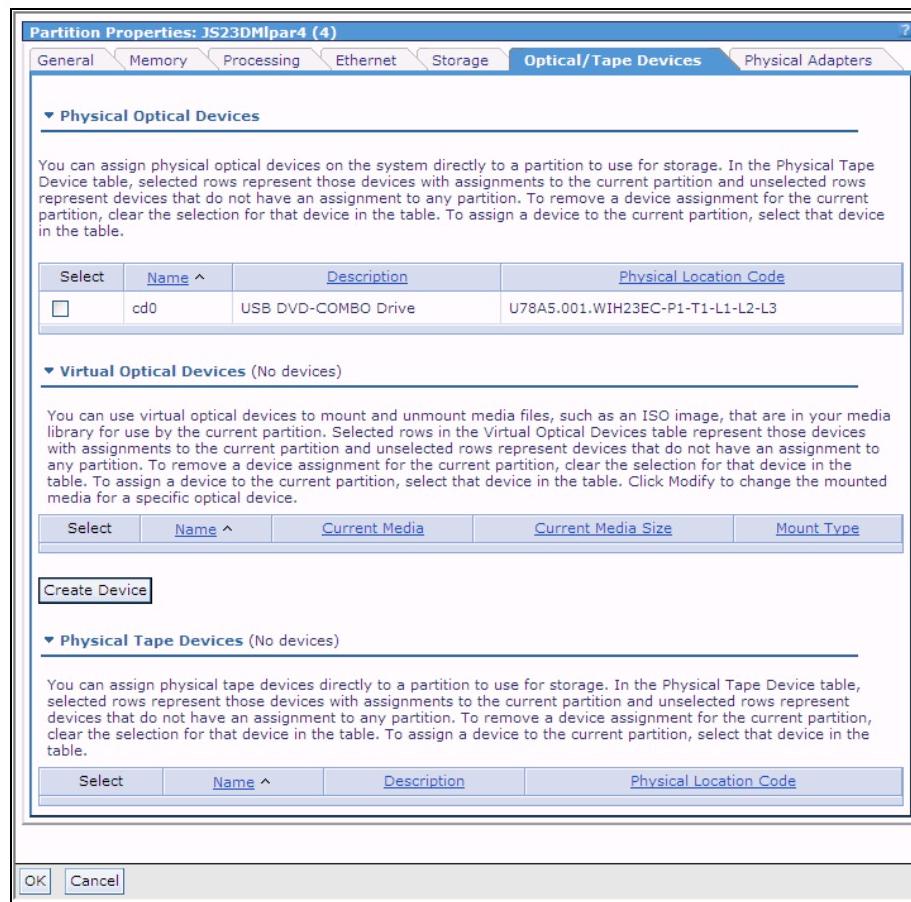


Figure 4-73 Partition Properties Optical/Tape Devices tab

Additional virtual optical devices can be created, and the media that is mounted to an existing virtual optical device can be changed in this window. Creating virtual optical media is covered in “Virtual optical devices” on page 135.

Figure 4-74 on page 165 shows an existing virtual optical device, vt0pt0, being modified by changing the current media.

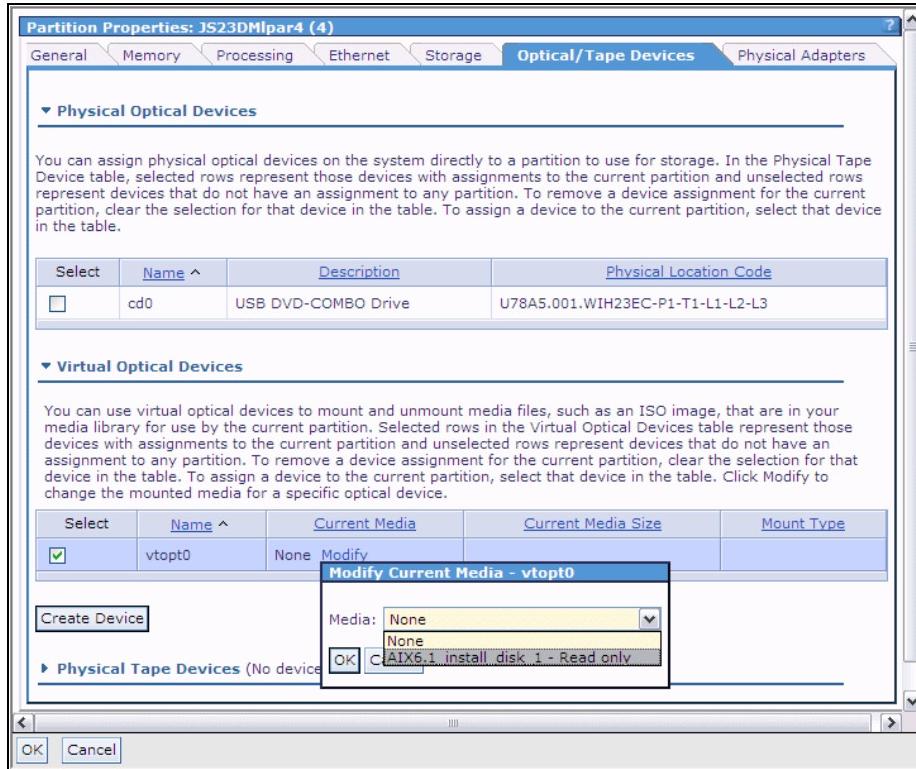


Figure 4-74 Partition Properties, changing the current mounted media

Physical adapters that are not assigned to an LPAR or any physical adapters that are already assigned to the selected LPAR will be displayed when the **Physical Adapters** tab is clicked.

Figure 4-75 on page 166 shows a Gigabit Ethernet-SX PCI-X Adapter available for assignment to this LPAR.

Note: Partitions that are configured for shared memory or IBM i partitions cannot use physical adapters. Therefore, the Physical Adapter tab of this window will not be shown when viewing the properties of these types of LPARs.

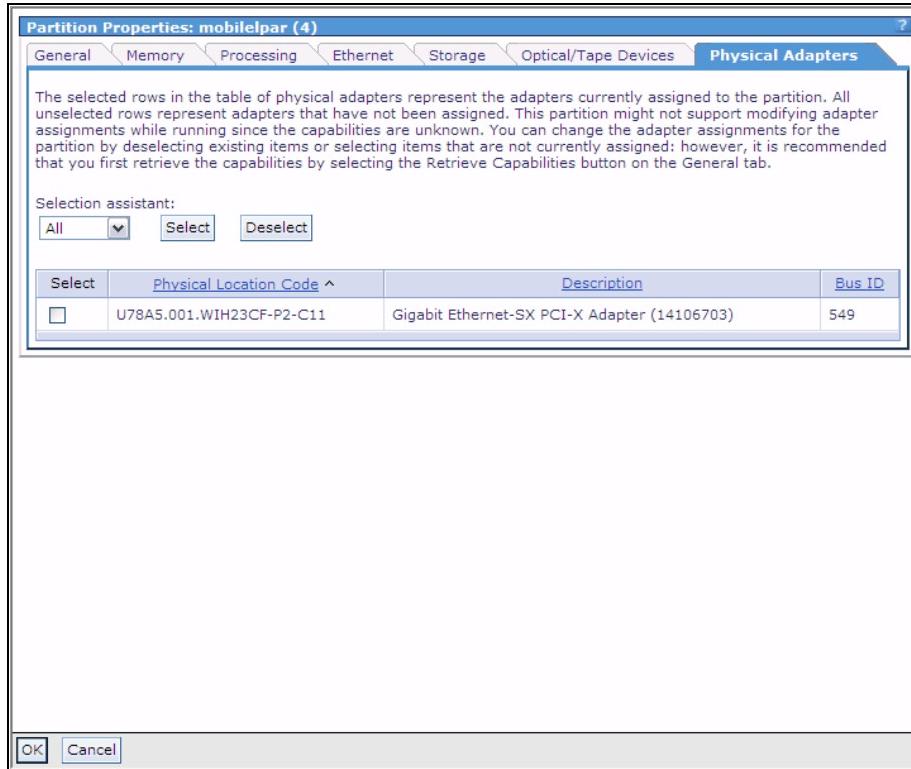


Figure 4-75 Partition Properties, Physical Adapters tab

Note: The Partition Properties window for the VIOS partition does not have the **Storage** and **Optical Devices** tabs.

4.8 Console access and activating a partition

The following sections discuss basic access to a partition and partition management functions.

4.8.1 Opening a virtual terminal

Accessing a partition virtual terminal from the VIOS can be done in two different ways. However, only one virtual terminal to an LPAR can be open at a time.

Note: These methods are not available for IBM i. In the case of IBM i, the Operations Console (LAN) is the only supported system console.

The first method from the IVM UI is shown in Figure 4-76 on page 167. From the View/Modify Partitions view, select the check box for the desired LPAR. From the drop-down box, select **Open terminal window**.

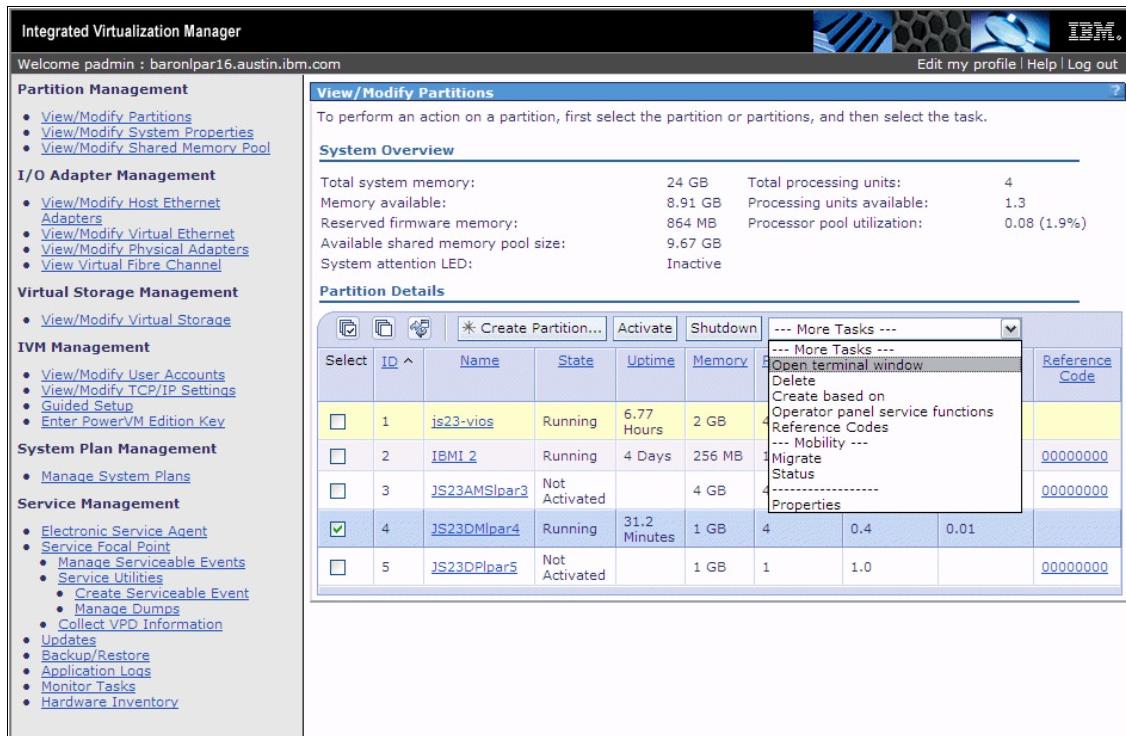


Figure 4-76 Start virtual terminal for an LPAR

A separate applet will start and open the virtual terminal window. This new window will prompt for the password of the login ID of the current IVM session.

Figure 4-77 shows a successful connection to the LPAR virtual terminal.

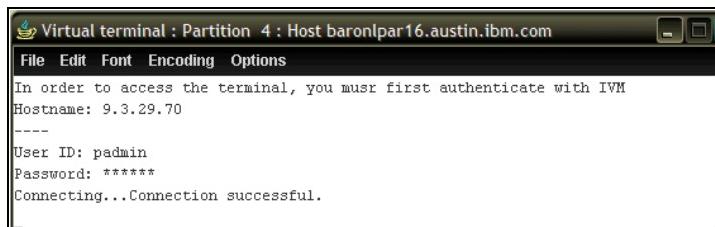


Figure 4-77 Virtual Terminal started from IVM UI

The second method to start a virtual terminal is from the VIOS command line. From the command line prompt, issue the **mkvt** command as shown in Example 4-23.

Example 4-23 Creating a virtual terminal from the command line

```
$ mkvt -id 4
```

Specify the partition number that you want to connect after the **-id** flag. Your command line session to the VIOS will now become the virtual terminal for the specified partition.

Note: The key sequence <cr>~. or enter key, tilde, period entered from the virtual terminal will allow you to break out of the **mkvt** command or close the virtual terminal applet.

4.8.2 Activating a partition

During initial setup and configuration of an LPAR, open a virtual terminal connection to the partition prior to activating.

Activating from the UI

Activating a partition from the IVM UI starts from View/Modify Partitions. In this view, select the partition to be activated and click the **Activate** button as shown in Figure 4-78.

Integrated Virtualization Manager

Welcome padmin : baronpar16.austin.ibm.com

Partition Management

- [View/Modify Partitions](#)
- [View/Modify System Properties](#)
- [View/Modify Shared Memory Pool](#)

I/O Adapter Management

- [View/Modify Host Ethernet Adapters](#)
- [View/Modify Virtual Ethernet](#)
- [View/Modify Physical Adapters](#)
- [View Virtual Fibre Channel](#)

Virtual Storage Management

- [View/Modify Virtual Storage](#)

IVM Management

- [View/Modify User Accounts](#)
- [View/Modify TCP/IP Settings](#)
- [Guided Setup](#)
- [Enter PowerVM Edition Key](#)

System Plan Management

- [Manage System Plans](#)

Service Management

- [Electronic Service Agent](#)
- [Service Focal Point](#)
 - [Manage Serviceable Events](#)
 - [Service Utilities](#)
 - [Create Serviceable Event](#)
 - [Manage Dumps](#)
 - [Collect VPD Information](#)
 - [Updates](#)
 - [Backup/Restore](#)
 - [Application Logs](#)
 - [Monitor Tasks](#)
 - [Hardware Inventory](#)

Figure 4-78 Activating a partition

The next window shows the current state of the partition and asks you to confirm activation by clicking **OK**, as shown in Figure 4-79 on page 170.

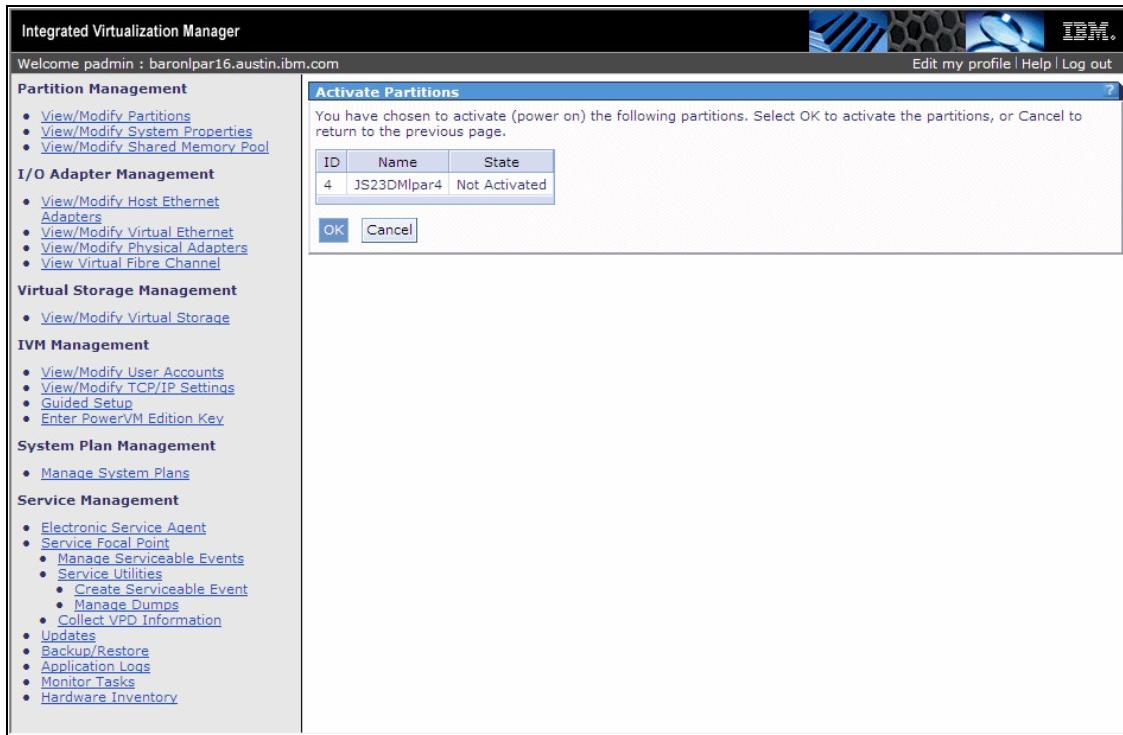


Figure 4-79 Confirm partition activation

When the LPAR activation starts, the message Performing Task - Please Wait will briefly appear, then the IVM UI will return to the View/Modify Partitions window.

Activating from the CLI

The **chsysstate** command is used to start a partition from the command line by either the LPAR number or name. Example 4-24 shows LPAR 4 being activated from the CLI.

Example 4-24 Using the CLI to activate a partition

```
$ chsysstate -o on -r lpar --id 4
```

or

```
$ chsysstate -o on -r lpar -n JS23DMlpar4
```

The **lsrefcode** command can be used to monitor the status codes as the LPR becomes active. Example 4-25 shows the lsrefcode being used with both LPAR number and name for LPAR 4.

Example 4-25 Using the CLI to monitor partition activation status codes

```
$ lsrefcode -r lpar --filter lpar_names=JS23DM1par4 -F refcode  
00000000  
$ lsrefcode -r lpar --filter lpar_names=JS23DM1par4 -F refcode  
CA00E1F1  
$ lsrefcode -r lpar --filter lpar_names=JS23DM1par4 -F refcode  
AA00E158
```

or

```
$ lsrefcode -r lpar --filter lpar_ids=4 -F refcode  
00000000  
$ lsrefcode -r lpar --filter lpar_ids=4 -F refcode  
CA00E1F1  
$ lsrefcode -r lpar --filter lpar_ids=4 -F refcode  
AA00E158
```

The **lssyscfg** command can be used to display the state of the LPARS by name or the ID number of the LPAR, as shown in Example 4-26.

Example 4-26 The lssyscfg command used to display the LPAR state

```
$ lssyscfg -r lpar -F name,state  
js23-vios,Running  
IBMI 2,Running  
JS23AMS1par3,Not Activated  
JS23DM1par4,Running  
JS23DP1par5,Not Activated
```

or

```
$ lssyscfg -r lpar -F lpar_id,state  
1,Running  
2,Running  
3,Not Activated  
4,Running  
5,Not Activated
```

4.8.3 Shutdown a VIO Client partition

The shutdown of a partition can be initiated from the UI or the CLI. The shutdown process can interact with the operating system on an LPAR, or can be immediate without notifying the operating system.

The following options are available for a partition shutdown

- ▶ Operating System (recommended)
- ▶ Delayed
- ▶ Immediate

The Operating System shutdown option is available only if the RMC connection is active. It is the recommended method. The Delayed option is the equivalent of pushing the white control-panel power button. AIX partitions will gracefully handle this option, but Linux partitions are required to install IBM service and productivity tools for Linux on POWER.

The Immediate shutdown option should be used only as a last resort because of the potential of data loss.

IBM i partitions are recommended to be shut down by using the IBM i session commands SHTDWNSYS, ENDSBS, or PWRDWNSYS.

Shutdown from the UI

From the View/Modify Partitions view, select the check box for the LPAR to be shut down and then click the **Shutdown** button as shown in Figure 4-80 on page 173.

Integrated Virtualization Manager

Welcome padmin : baronpar16.austin.ibm.com

Partition Management

- [View/Modify Partitions](#)
- [View/Modify System Properties](#)
- [View/Modify Shared Memory Pool](#)

I/O Adapter Management

- [View/Modify Host Ethernet Adapters](#)
- [View/Modify Virtual Ethernet](#)
- [View/Modify Physical Adapters](#)
- [View Virtual Fibre Channel](#)

Virtual Storage Management

- [View/Modify Virtual Storage](#)

IVM Management

- [View/Modify User Accounts](#)
- [View/Modify TCP/IP Settings](#)
- [Guided Setup](#)
- [Enter PowerVM Edition Key](#)

System Plan Management

- [Manage System Plans](#)

Service Management

- [Electronic Service Agent](#)
- [Service Focal Point](#)
 - [Manage Serviceable Events](#)
 - [Service Utilities](#)
 - [Create Serviceable Event](#)
 - [Manage Dumps](#)
 - [Collect VPD Information](#)
 - [Updates](#)
 - [Backup/Restore](#)
 - [Application Logs](#)
 - [Monitor Tasks](#)
 - [Hardware Inventory](#)

Figure 4-80 Shutdown an LPAR

The Shutdown partitions window, as shown in Figure 4-81 on page 174, will be displayed.

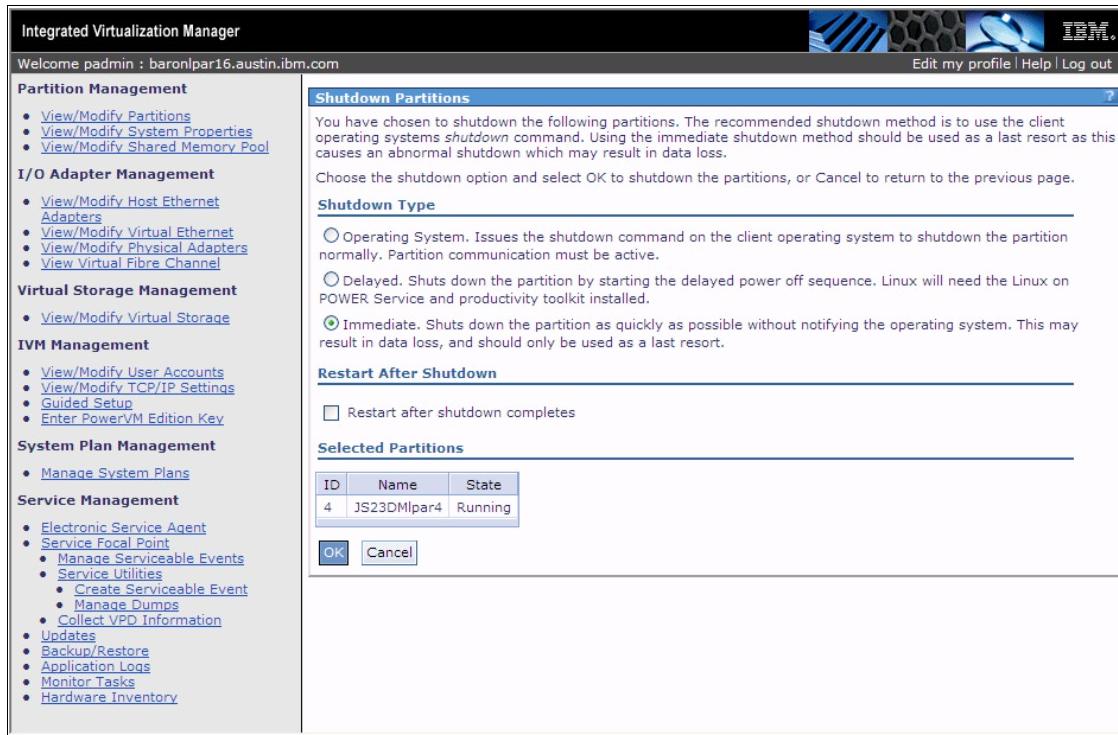


Figure 4-81 Partition shutdown options

From this window, choose the shutdown type option. The partition can also be restarted after the shutdown by selecting the restart check box option. Click **OK** and the partition will be shut down.

Note: The Operating System option will be disabled if RMC is not active between the LPAR and VIOS. The Delayed option will be selected by default.

Shutdown from the CLI

The **chsysstate** command is used to shut down a partition from the command line by either the LPAR number or name, as shown in Example 4-27. This example uses the operating system shutdown option on LPAR 4.

Example 4-27 CLI shutdown of a logical partition

```
$ chsysstate -r lpar -o osshutdown --id 4
```

or

```
$ chsysstate -r lpar -o osshutdown -n "JS23DM1par4"
```

The corresponding CLI shutdown options to use with the **-o** flag are:

- ▶ `osshutdown` (Operating System)
- ▶ `shutdown` (Delayed, white button shutdown)
- ▶ `shutdown --immed` (Immediate)

4.8.4 Shutdown the VIO Server

The VIO Server is shutdown in a similar process to a VIO Client LPAR. Both the UI and CLI can be used.

Shutdown from the UI

When selecting the VIOS partition to be shut down, a warning is presented stating that shutting down the IVM partition will shut down all partitions and the entire system. There is no shutdown option to choose only the restart option.

Figure 4-82 on page 175 shows the VIOS Shutdown Partitions window.

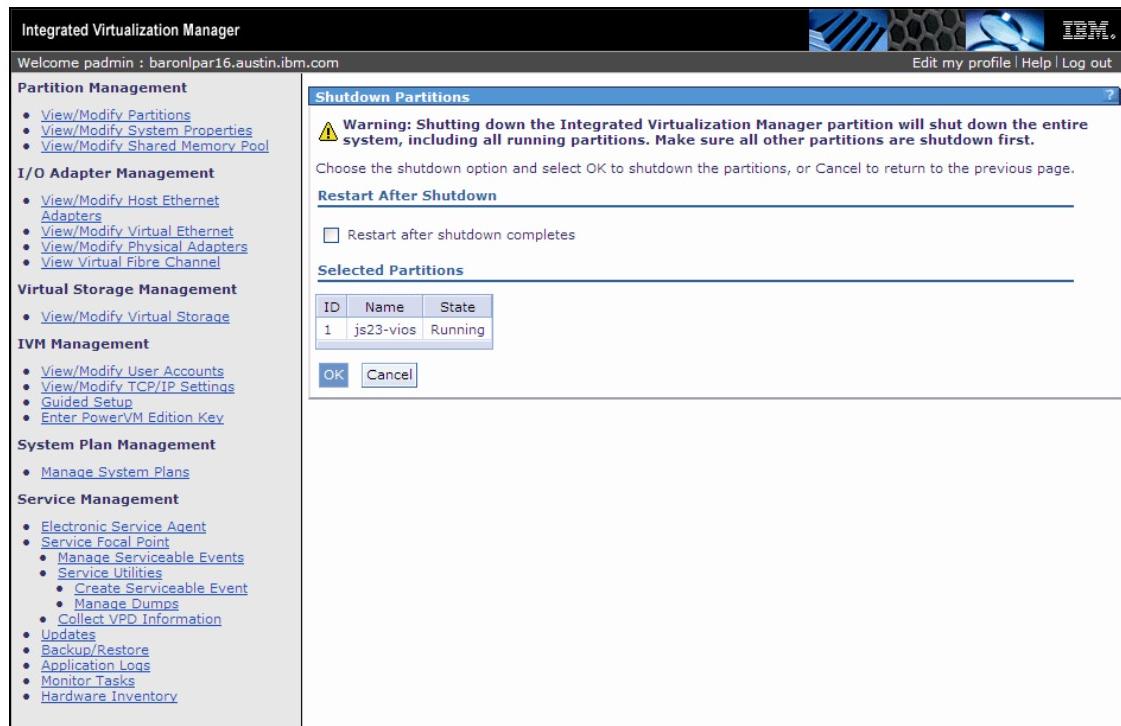


Figure 4-82 VIOS partition shutdown options

Shutdown using the CLI

The **shutdown** command to use from the CLI or console session is shown in Example 4-28 and has two options. To automatically restart after the shutdown use the **-restart** flag, to suppress the warning message and confirmation add the **-force** option.

Example 4-28 VIOS shutdown command

```
$ shutdown -restart  
Shutting down the VIO Server could affect Client Partitions. Continue  
[y|n]?
```

or

```
$ shutdown -restart -force
```



Active Memory Sharing configuration using IVM

Active Memory Sharing is an IBM PowerVM advanced memory virtualization technology that provides system memory virtualization capabilities to IBM Power Systems, allowing multiple logical partitions to share a common pool of physical memory.

This chapter describes how to configure Active Memory Sharing (AMS) using the IVM UI and at a high level some of the planning considerations that should be used. Additional information on Active Memory Sharing can be found at the IBM PowerVM portal located here:

<http://www-03.ibm.com/systems/power/software/virtualization/>

Also *PowerVM Virtualization Active Memory Sharing*, REDP-4470, as well as *IBM PowerVM Active Memory Sharing: An Overview* Whitepaper (<ftp://ftp.software.ibm.com/common/ssi/sa/wh/n/pow03026usen/POW03026USEN.PDF>) publications include a more complete technical discussion of Active Memory Sharing.

This chapter has the following sections:

- ▶ “Planning considerations” on page 179
- ▶ “Configuration of Active Memory Sharing using IVM” on page 183

- ▶ “Active Memory Sharing summary” on page 209

5.1 Planning considerations

Active Memory Sharing is an IBM PowerVM advanced memory virtualization technology that provides system memory virtualization capabilities to IBM Power Systems, allowing multiple logical partitions to share a common pool of physical memory.

When using a shared memory mode, it is the system that automatically decides the optimal distribution of the physical memory to logical partitions and adjusts the memory assignment based on demand for memory pages. The administrator just reserves physical memory for the shared memory pool and assigns logical partitions to the pool."

"The PowerVM™ Active Memory™ Sharing technology (or shared memory) allows logical partitions to share the memory in the shared memory pool. Instead of assigning a dedicated amount of physical memory to each logical partition that uses shared memory (hereafter referred to as shared memory partitions), the hypervisor constantly provisions the physical memory from the shared memory pool to the shared memory partitions as needed. The hypervisor allows shared memory partitions that need memory to use portions of the shared memory pool that are not currently being used by other shared memory partitions. When a shared memory partition needs more memory than the current amount of unused memory in the shared memory pool, the hypervisor stores a portion of the memory that belongs to the shared memory partition in auxiliary storage. Access to the auxiliary storage is provided by a Virtual I/O Server logical partition. When the operating system attempts to access data that is located in the auxiliary storage, the hypervisor directs a Virtual I/O Server to retrieve the data from the auxiliary storage and write it to the shared memory pool so that the operating system can access the data. The PowerVM Active Memory Sharing technology is available with the PowerVM Enterprise Edition hardware feature, which also includes the license for the Virtual I/O Server software.

Successful implementation of Active Memory Sharing requires a thorough understanding of the workloads that will be placed on memory sharing partitions. Workloads that have a high sustained memory needs because of sustained peak loads or High Performance Computing applications would not be good candidates for Active Memory Sharing.

5.1.1 Prerequisites

To use Active Memory Sharing on a IBM BladeCenter JS23 or JS43 minimum software, configuration and feature levels must be met. These are:

- ▶ IBM PowerVM Enterprise Edition

- ▶ Virtual Input/Output Server 2.1.1
- ▶ Only virtual I/O, no physical adapters or logical ports from an HEA allowed
- ▶ Only shared processor mode, no dedicated processors
- ▶ AIX 6.1 TL3
- ▶ IBM i 6.1 plus latest cumulative PTF package + SI32798
- ▶ SUSE Linux Enterprise Server 11

5.1.2 Workload

In the Active Memory Sharing environment there are three types of shared memory scenarios:

- ▶ Physical over-commit
- ▶ Logical over-commit
- ▶ Non over-commit

Physical over-commit occurs when the combined working set of memory, of all the shared memory partitions, is greater than the total physical memory in the shared memory pool. In this case the actively used logical memory of the shared memory is backed by physical memory and the paging devices.

Logical over-commit occurs when actively referenced memory pages by the partitions are backed by physical memory allocated from the Active Memory Sharing Shared Memory Pool. Logical memory that is not actively referenced and that exceeds the amount of physical memory will reside on the paging device.

Non-overcommit occurs when the backing physical memory in the shared memory pool is adequate to cover all the peak needs of the shared memory partitions.

The potential workload candidates for Active Memory Sharing should be monitored for a period of time using a dedicated memory partition to understand the size and range requirements of memory usage.

Physical over-commit would be suitable for workloads such as file and print servers and network applications that are not highly sensitive to I/O latency. Shared memory partitions in this model will require paging of active memory pages.

Logical over-commit scenarios would work best with partition workloads that have peak usage at different times of the day causing utilization peaks and valleys to overlap. Other examples that would be prime candidates are test and

development environments and workloads that do not have sustained load requirements.

5.1.3 Paging devices

Active Memory Sharing paging devices and operating systems paging device considerations are similar. Active Memory Sharing paging operations will be typically be 4k in size. Write and read caches should be enabled. Striped disk configurations should be used when possible with a 4k stripe size.

The sizes of the dedicated paging devices for AIX and Linux shared memory partitions must at least equal the maximum amount of logical memory assigned. IBM i shared memory partitions must have paging devices that are at least the size of the maximum logical memory assigned multiplied by 1.008.

5.1.4 Collaborative Memory Manager

Collaborative Memory Manager (CMM) running in the operating system of the shared memory partitions coordinates with the hypervisor to provide hints on logical page usage. These hints of the status of memory pages will help the hypervisor prioritize pages.

CMM is configured from the OS running in the shared memory partition. When the CMM loan policy is enabled, the OS loans pages to the hypervisor to help reduce hypervisor paging required to support the overcommitment of memory.. When CMM policy is disabled the hypervisor will do the paging based on hints from the OS.

5.1.5 CPU resource sizing

This section looks at VIO Server and shared memory partition additional CPU entitlements required in an Active Memory Sharing environment.

Virtual I/O Server

Additional CPU resources will be required for the VIO Server to properly service the Active Memory Sharing environment. The hypervisor does the address translation and allocation of physical memory across partitions and the VIOS donates cycles to the hypervisor for these operations and also needs cycles for the paging activity.

The additional VIO Server CPU requirements per shared memory partition can typically range from 0.005 for light paging rates using internal storage to 0.16 for heavy paging rates using high end SAN storage as shown in Table 5-1.

Table 5-1 Estimated additional VIOS CPU entitlement per shared memory LPAR

Paging rate	Storage types			
	Internal storage	Entry level storage	Mid range storage	High end storage
Light	0.005	0.01	0.02	0.02
Moderate	0.01	0.02	0.04	0.08
Heavy	0.02	0.04	0.08	0.16

Shared memory partition

Shared memory partitions will also require additional CPU entitlement compared to dedicated memory partitions running the same workload. The additional amount will be dependent on the memory subscription ration (sum of logical memory divided by physical memory in the pool) and CMM loan policy settings. In case where memory is not over committed the additional CPU entitlement will be minimal.

IBM Workload Estimator (IBM WLE) will be updated to produce estimated sizing for shared memory partition CPU requirements.

IBM Workload Estimator can be found at:

<http://www-947.ibm.com/systems/support/tools/estimator/>

5.1.6 Memory weight

The hypervisor considers several variables to allocate and manage memory across the shared memory partitions. These variables included global and partition load, global and partition memory pressure, and the one direct user adjustment memory weight. Similar workloads that run concurrently require different weights if one partition requires a priority. In an active/inactive scenario when one shared memory partition uses most of the resources the memory weight should be equal in both partitions to enhance memory transfer when the inactive partition becomes active.

5.1.7 Consolidation factors

Once the workload types in a environment have been determined and selected for consolidation additional factors must be considered:

- ▶ The logical to physical subscription ratio must be suitable for the selected workloads.

- ▶ Assigning a memory weight. The IVM UI allows three values, low, medium and high, with a default of medium.
- ▶ Paging device configuration, the higher the subscription ratio the higher the need for optimized paging devices.
- ▶ CMM configuration determines page loaning policy. Application load and loaning policy, none to aggressive, should be evaluated for acceptable performance. CMM is set at the OS level therefore a mix of loaning levels can exist in the same system.
- ▶ Rebalance resources to improve overall system utilization. Previous memory limits that caused under utilization of CPU and memory resources can be reviewed for new allocations or additional workloads.

5.2 Configuration of Active Memory Sharing using IVM

IVM is used to configure the main components of Active Memory Sharing, the shared memory pool and the paging storage devices used by the hypervisor. Once Active Memory Sharing is configured IVM is used to create new, or change existing LPARS to utilize shared memory

5.2.1 Defining the shared memory pool and paging storage pool

When creating the shared memory pool for Active Memory Sharing you also create the paging storage pool in the same step. The shared memory pool cannot exceed the current memory available minus the amount required for Reserved Firmware Memory.

Note: Reserved Firmware Memory requirements increase when the Active Memory Sharing shared memory pool is created. Also as additional partitions are created and additional virtual resources are defined the Reserved Firmware Memory value will increase. If the Active Memory Sharing shared memory pool is defined close to the maximum amount you limit the ability to create additional resources.

There are two possibilities for the paging space that is managed by the hypervisor. The first is to create a common paging storage pool and this step is required. The second is to create individual paging devices for each LPAR using shared memory and is covered in 5.2.2, “Creating dedicated paging devices for partitions” on page 191

Since a common paging storage pool is required, the first step is to created a *storage pool* that can be assigned as the paging storage pool.

1. To create a common paging storage pool start in the navigation area of the IVM UI and click on **View/Modify Virtual Storage** as shown in Figure 5-1.

ID	Name	State	Uptime	Memory	Processors	Entitled Processing Units	Utilized Processing Units	Reference Code
1	is23-vios	Running	1.13 Days	2 GB	4	0.4	0.01	
2	IBMI_2	Not Activated		256 MB	1	0.5		00000000

Figure 5-1 Start Active Memory Sharing configuration with View/Modify Virtual Storage

2. The next window will begin the storage pool creation process. Select **Create Storage Pool** as shown in Figure 5-2 on page 185. This pool will then be used for shared paging storage.

Integrated Virtualization Manager

Welcome padmin : baronpar16.austin.ibm.com

Partition Management

- [View/Modify Partitions](#)
- [View/Modify System Properties](#)
- [View/Modify Shared Memory Pool](#)

I/O Adapter Management

- [View/Modify Host Ethernet Adapters](#)
- [View/Modify Virtual Ethernet](#)
- [View/Modify Physical Adapters](#)
- [View Virtual Fibre Channel](#)

Virtual Storage Management

- [View/Modify Virtual Storage](#)

IVM Management

- [View/Modify User Accounts](#)
- [View/Modify TCP/IP Settings](#)
- [Guided Setup](#)
- [Enter PowerVM Edition Key](#)

System Plan Management

- [Manage System Plans](#)

Service Management

- [Electronic Service Agent](#)
- [Service Focal Point](#)
 - [Manage Serviceable Events](#)
 - [Service Utilities](#)
 - [Create Serviceable Event](#)
 - [Manage Dumps](#)
 - [Collect VPD Information](#)
 - [Updates](#)
 - [Backup/Restore](#)
 - [Application Logs](#)
 - [Monitor Tasks](#)
 - [Hardware Inventory](#)

Figure 5-2 Starting the storage pool creation

- The next window will prompt for the name of the storage pool and the name must be a valid name for volume groups, for example no spaces are allowed and the name cannot exceed 15 characters. Use the choice of **Logical volume based** for the storage pool type. Next select the physical volume or volumes desired to create the pool as shown in Figure 5-3 on page 186.
- When the entries are made click the **OK** button to complete the storage pool creation process.

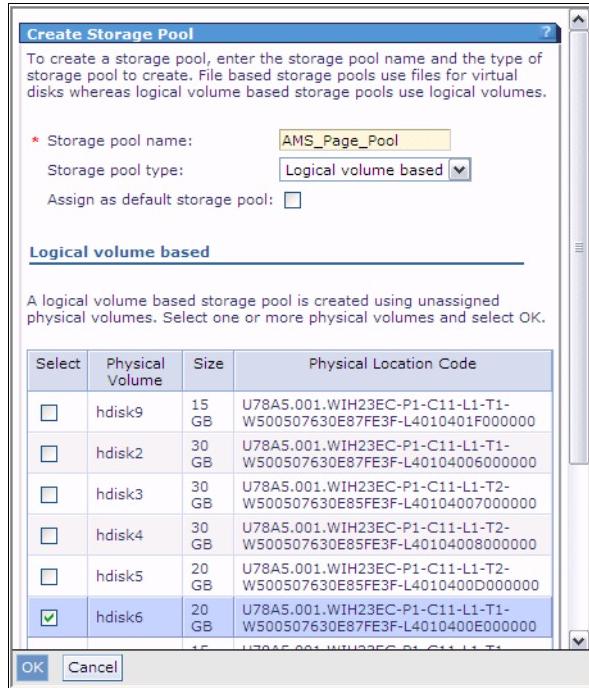


Figure 5-3 Naming the storage pool and selecting backing devices

Figure 5-4 on page 187 shows the newly created storage pool.

Select	Name	Total Size	Available Size	Type
<input type="checkbox"/>	AMS Page Pool	19.92 GB	19.92 GB	Logical volume based
<input type="checkbox"/>	media lib pool (Default)	29.91 GB	29.91 GB	Logical volume based
<input type="checkbox"/>	rootvg	68.25 GB	38 GB	Logical volume based

Figure 5-4 Storage pool list with new pool for paging added

Note: A new designation of “Paging” will be added to the name field of the storage pool list when the shared memory pool is created.

With the paging storage pool created we are ready to defined the shared memory pool. From the IVM UI click **View/Modify Shared Memory Pool**.

Figure 5-5 on page 188 shows the shared memory pool configuration page. The first items to note are the current memory available and the reserved firmware memory values. The amount of memory to be configured for the shared memory pool cannot exceed the difference between these two values.

Note: Reserved firmware memory will increased as additional LPARs, adapters, etc. are configured. If the shared memory pool is configured near to the maximum allowed value you may have to reduce the pool size in the future to create additional LPARs or create/additional adapters.

The screenshot shows the IVM interface with a sidebar containing various management categories like Partition Management, I/O Adapter Management, and Service Management. The main window is titled 'View/Modify System Properties' and is focused on the 'Memory' tab. It displays system memory details: Installed system memory (24 GB), Configurable system memory (24 GB), Current memory available (21.34 GB), Pending memory available (21.34 GB), and Reserved firmware memory (416 MB). Below this, there's a section for 'Memory Region Size' where the current value is 32 MB, and a dropdown menu also shows 32 MB. A note indicates that a shared memory pool defines the amount of shared memory available on the system. At the bottom are 'Apply' and 'Reset' buttons.

Figure 5-5 Defining a shared memory pool

- Clicking the **Define Shared Memory Pool** button will open the dialog for input of the desired memory pool size and the storage pool to be used for the paging storage pool. When these values have been entered and selected from the drop down box as shown in Figure 5-6 on page 189 click the **OK** button.

Note: When IVM creates the shared memory pool, the value provided for the **Assigned memory** of the pool will also be used for the maximum value of the pool.

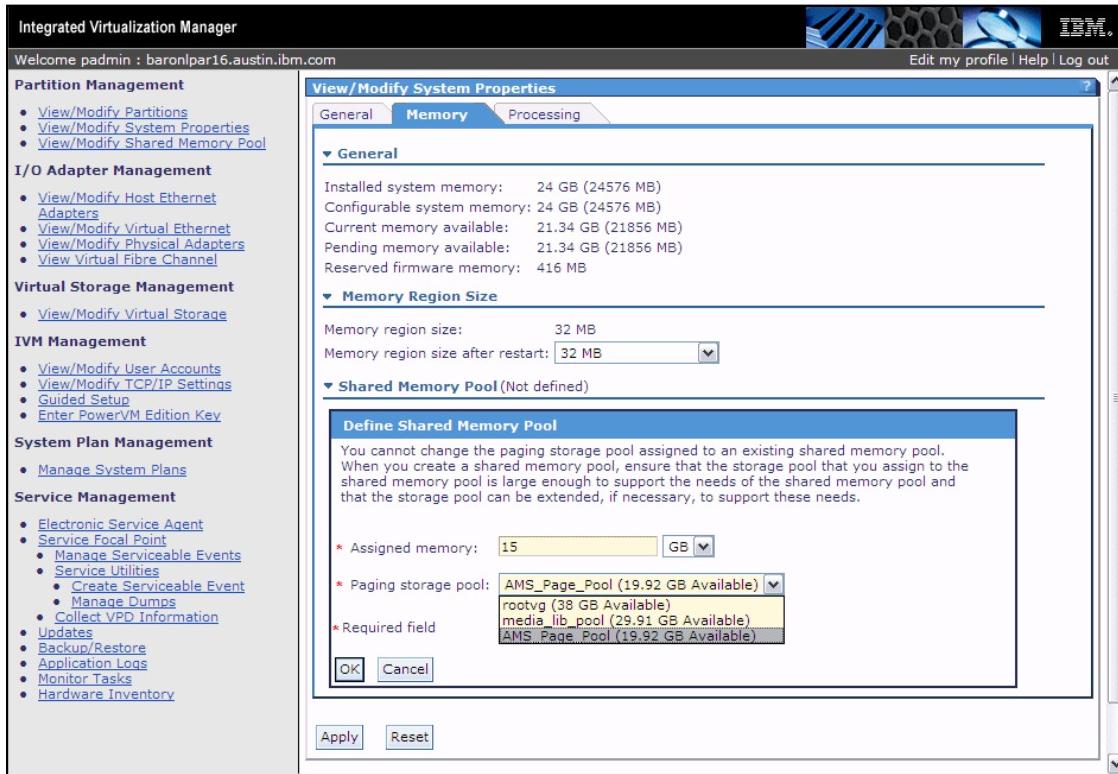


Figure 5-6 Shared memory pool configuration values

5. After clicking the **OK** button the screen will refresh and indicate the shared memory pool has been defined as shown in Figure 5-7 on page 190.

Figure 5-7 Shared memory pool defined state

6. Click the **Apply** button to create the shared memory pool and the assignment of the paging storage pool as shown in Figure 5-8 on page 191.

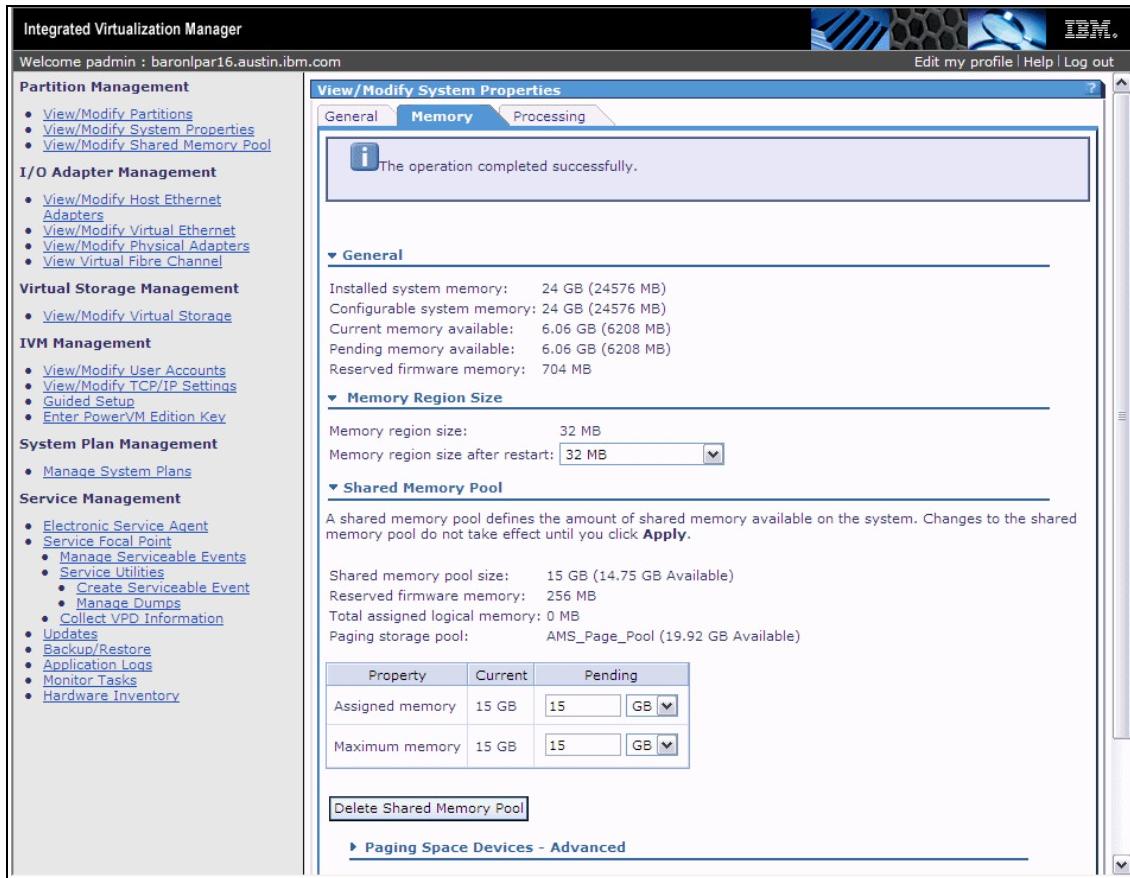


Figure 5-8 Shared memory pool information after creation

Now that we have created a shared memory pool we can create LPARs that use shared memory. As these LPARs are created, Active Memory Sharing will subdivide the paging storage pool through the use of logical volumes to accommodate each LPAR.

The recommended method however is to provide dedicated physical devices for each LPAR using shared memory as hypervisor paging devices. The next section will detail how these dedicated paging devices are created.

5.2.2 Creating dedicated paging devices for partitions

To create dedicated physical devices for LPAR paging space perform the following steps:

1. Click the **View/Modify Shared Memory Pool** from the IVM navigation area.
2. Then click on **Paging Space Devices - Advanced** to expand the section as shown in Figure 5-9 on page 192.

The screenshot shows the IVM interface with the following details:

- Navigation Area:**
 - Integrated Virtualization Manager
 - Welcome padmin : baronlpars16.austin.ibm.com
 - IBM logo and navigation links: Edit my profile | Help | Log out
- Left Sidebar (Partition Management):**
 - View/Modify Partitions
 - View/Modify System Properties
 - View/Modify Shared Memory Pool
- I/O Adapter Management:**
 - View/Modify Host Ethernet Adapters
 - View/Modify Virtual Ethernet
 - View/Modify Physical Adapters
 - View Virtual Fibre Channel
- Virtual Storage Management:**
 - View/Modify Virtual Storage
- IVM Management:**
 - View/Modify User Accounts
 - View/Modify TCP/IP Settings
 - Guided Setup
 - Enter PowerVM Edition Key
- System Plan Management:**
 - Manage System Plans
- Service Management:**
 - Electronic Service Agent
 - Service Focal Point
 - Manage Serviceable Events
 - Service Utilities
 - Create Serviceable Event
 - Manage Dumps
 - Collect VPD Information
 - Updates
 - Backup/Restore
 - Application Logs
 - Monitor Tasks
 - Hardware Inventory

A shared memory pool defines the amount of shared memory available on the system. Changes to the shared memory pool do not take effect until you click **Apply**.

Shared memory pool size: 15 GB (14.75 GB Available)

Reserved firmware memory: 256 MB

Total assigned logical memory: 0 MB

Paging storage pool: AMS_Page_Pool (19.92 GB Available)

Property	Current	Pending
Assigned memory	15 GB	<input type="text" value="15"/> <input type="button" value="GB"/>
Maximum memory	15 GB	<input type="text" value="15"/> <input type="button" value="GB"/>

[Delete Shared Memory Pool](#)

Paging Space Devices - Advanced

A paging space device is a block storage device that is dedicated to the shared memory pool. When assigned to a shared memory partition, the paging space device provides paging space for the partition, as needed. When you create or modify a shared memory partition, IVM creates and manages the required paging space device for the partition automatically. However, you can define a specific paging space device for the shared memory pool, such as a physical volume. IVM can then assign the paging space device to a partition when you create it, if the device meets the appropriate requirements.

Click Add to define a new paging space device for the shared memory pool, or select a device and click Remove.

Add...	Remove	Select	Name ^	Storage Pool	Assigned Partition	Partition State	Size

Figure 5-9 Creating dedicated paging devices for LPARS using shared memory

3. The **Add** button is clicked next to display the devices that are available for selection. Choose the desired physical volume device and click the **OK** button as shown in Figure 5-10 on page 193

Note: In the list of available devices virtual disks, if available, will also be displayed with physical volumes.

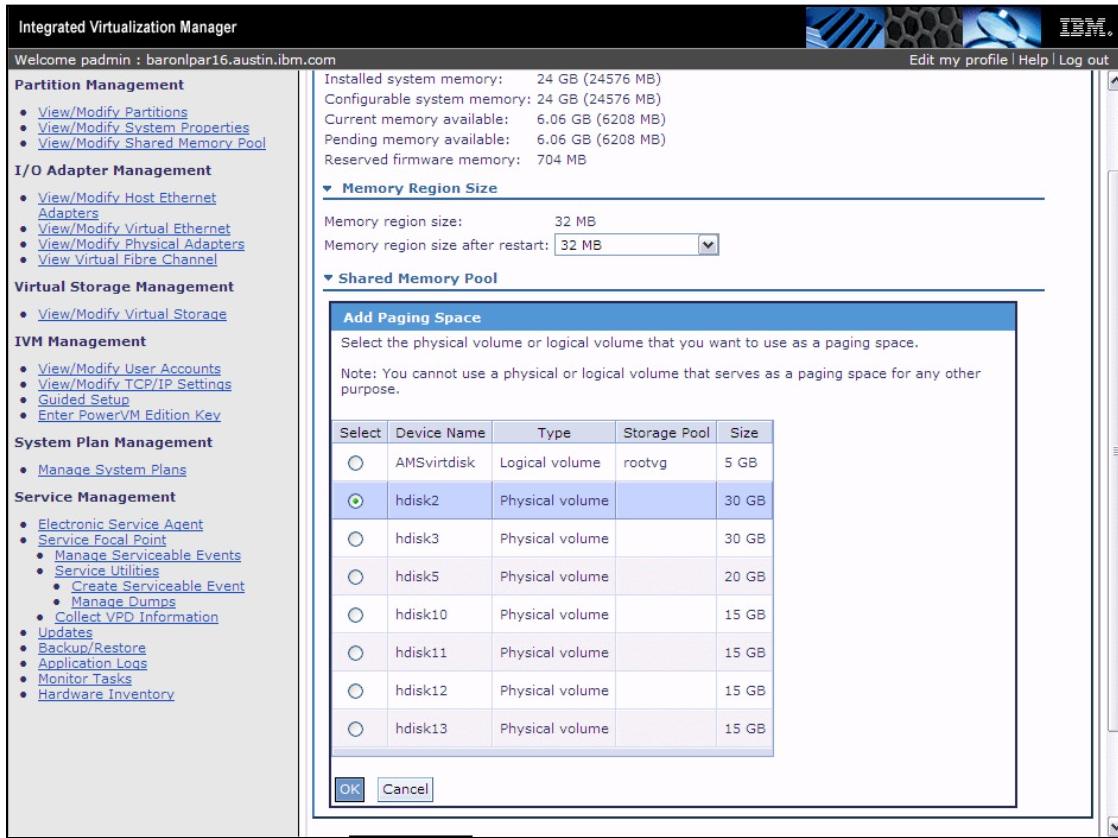


Figure 5-10 Dedicated device selection for share memory use

4. Figure 5-11 on page 194 shows the selected device now defined as a paging device. The **Apply** button must be clicked to complete the process.

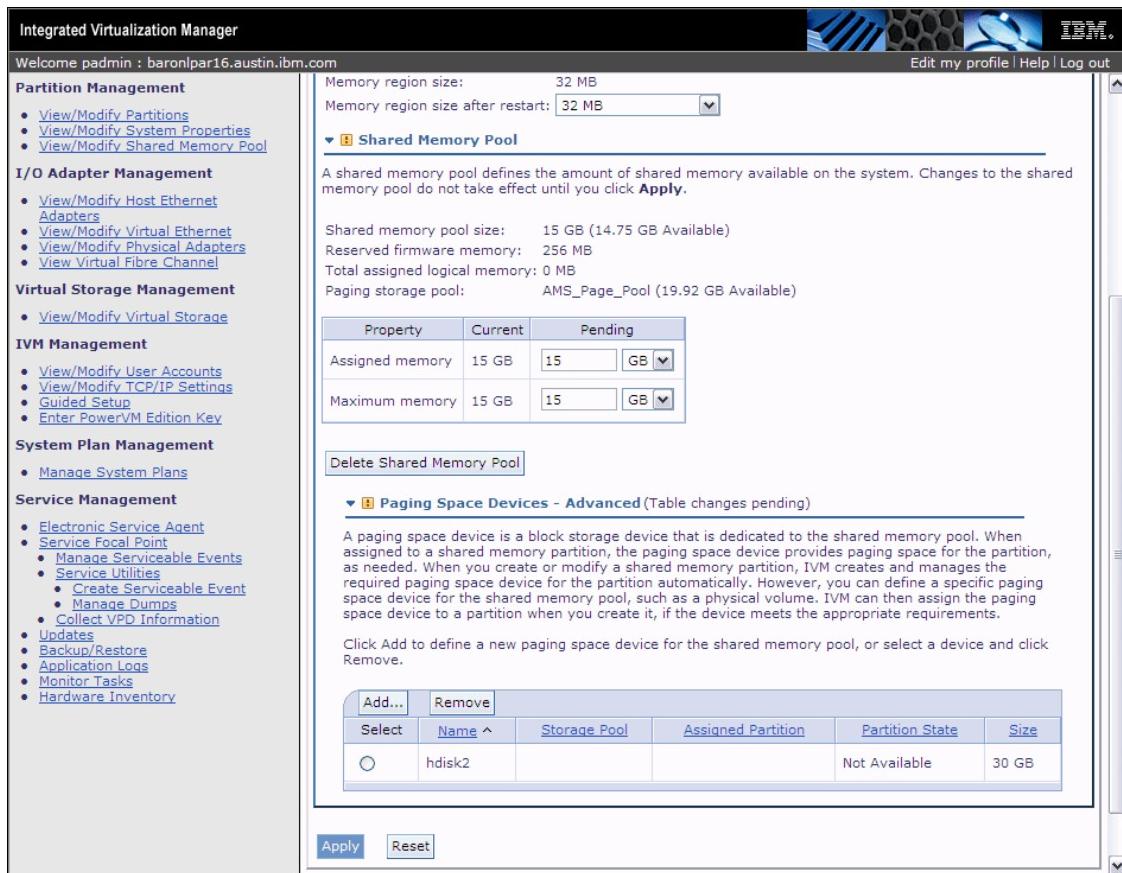


Figure 5-11 Dedicated device defined to paging space devices

Note: As LPARs are created that use shared memory, they will be assigned to the smallest dedicated device available that will meet memory size requirement.

5.2.3 Creating shared memory LPARs

Creating LPARs that used shared memory instead of dedicated memory uses the same wizard and process as detailed in 4.4.2, “Verify and modify VIOS partition memory and processors” on page 95.

However some of the options that were not selectable prior to the creation of a shared memory pool can now be chosen. Also options such as selecting logical

ports from a HEA, dedicated processors, and physical adapters are no longer available.

1. The LPAR wizard is started by clicking the **View/Modify Partitions** link on the IVM UI, and then clicking the **Create Partition** button. Figure 5-12 shows the first window of the wizard where the partition ID, partition name, and operating system environment are set. Enter the required information and click **Next**.

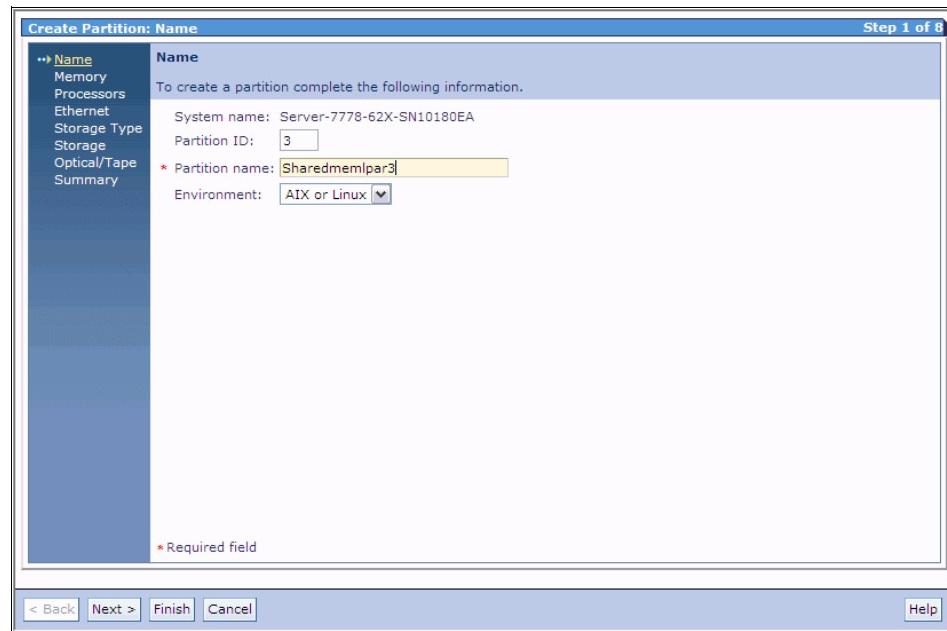


Figure 5-12 Creating a shared memory partition name

2. In the partition memory selection window shown in Figure 5-13 on page 196 the **Shared** option is now available and should be selected. The screen will refresh with shared mode memory information and an area to enter the desired amount of memory. The information display will show the total amount of shared memory in the pool and the cumulative total of assigned logical memory. The first shared memory LPAR created will indicate 0MB in this field. Enter the desired amount of logical memory in the box and click **Next**.

Note: When creating a partition in shared or dedicated mode, using the **Create Partition** wizard, the memory amount entered will be used for the assigned and maximum values.

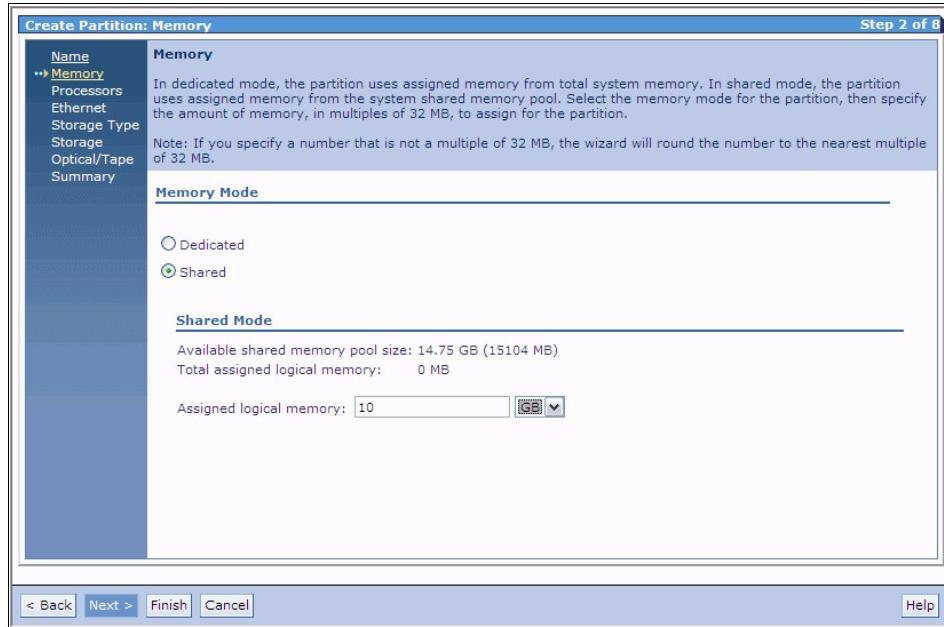


Figure 5-13 Selecting memory mode and amount for a shared memory partition

3. The next step will be the selection of the number of shared (virtual) processors as shown in Figure 5-14 on page 197. Notice that the dedicated processor option cannot be selected. Use the drop down box to select the number of assigned processors and click the **Next** button.

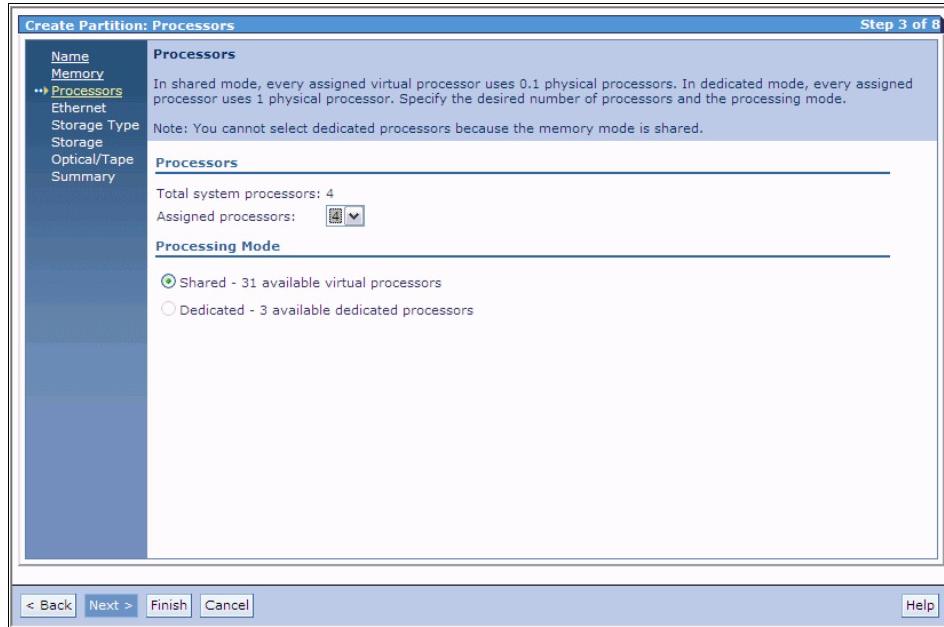


Figure 5-14 Selecting the number of processors in a shared memory partition

4. The next configuration step is Ethernet selection. As shown in Figure 5-15 on page 198 the only options are virtual Ethernet adapters. In this example we are using an existing Share Ethernet Adapter (SEA). Click on the **Next** button to continue to the storage options.

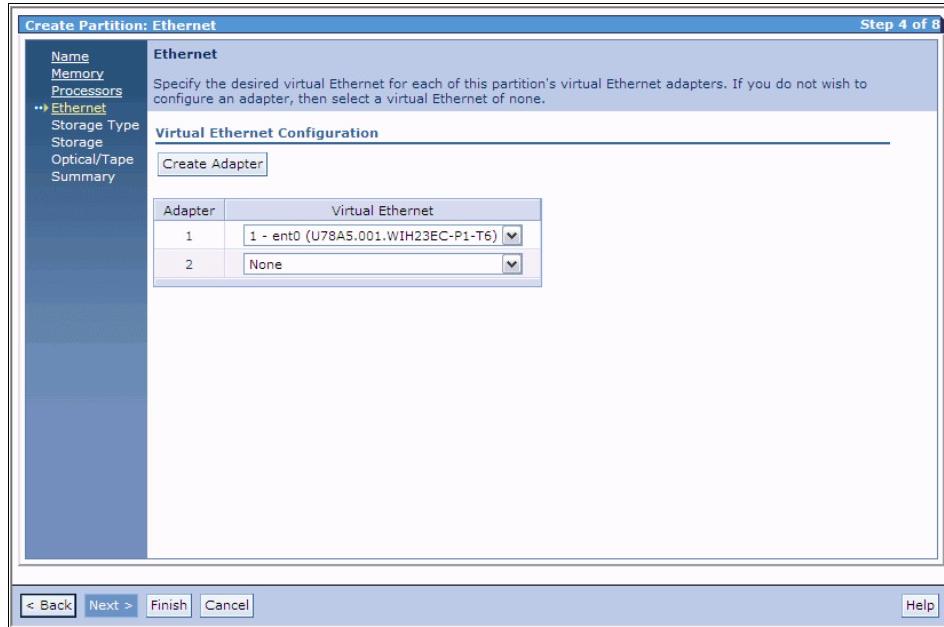


Figure 5-15 Ethernet selection for a shared memory partition

The storage selection options for a shared memory LPAR are the same as a dedicated memory LPAR. Virtual disks can be created from a existing storage pool. Existing virtual disks or physical volumes can be selected. There is also the None option if you do not desire to assigned storage at this time.

5. In Figure 5-16 on page 199 we chose the **Assign existing virtual disks and physical volumes** option. Click the **Next** button to continue to the selection window.

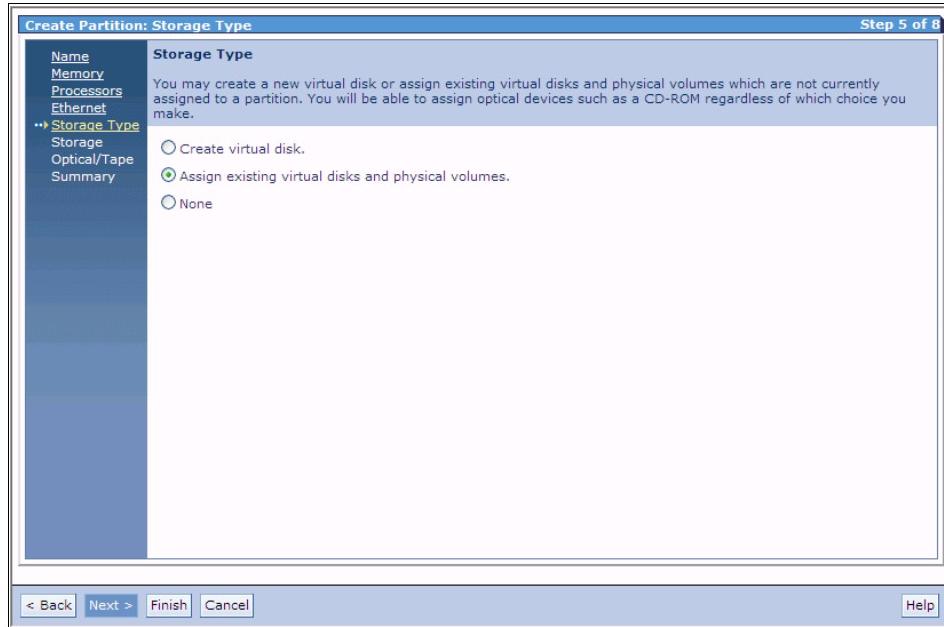


Figure 5-16 Storage selection for a shared memory partition

Figure 5-17 on page 200 shows the available selection of virtual disks (none in this example) and physical volumes that have not been assigned and are available.

6. In this example we chose hdisk4. Click the **Next** button to continue to the optical and tape options.

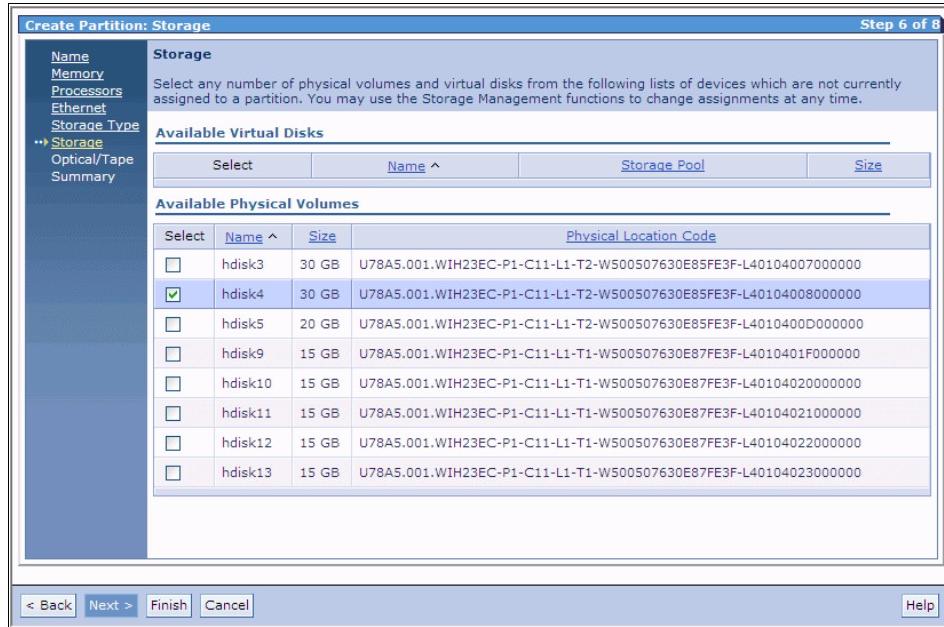


Figure 5-17 Storage selection for a shared memory partition

Two of the options shown in Figure 5-18 on page 201, physical optical devices and physical tape devices, will virtualize the physical device to the LPAR through the VIOS. Selecting these options does not imply a direct physical connection from the LPAR being created to the device. The virtual optical device is selected by default and can have media from the virtual media library assigned at this time.

7. Click the **Next** button to continue to the summary page.

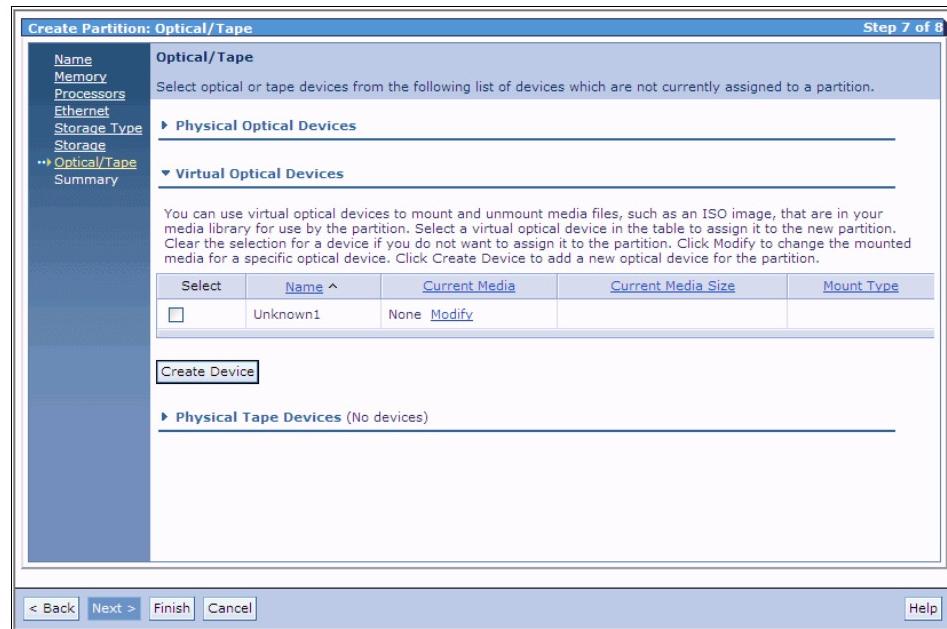


Figure 5-18 Optical and tape selections for a shared memory partition

The summary page as shown in Figure 5-19 on page 202 lists all of the selections made when stepping through the Create partition wizard.

8. The **Back** button can be used to revise any choices. Once the selections have been reviewed click the **Finish** button to complete the creation of the shared memory partition.

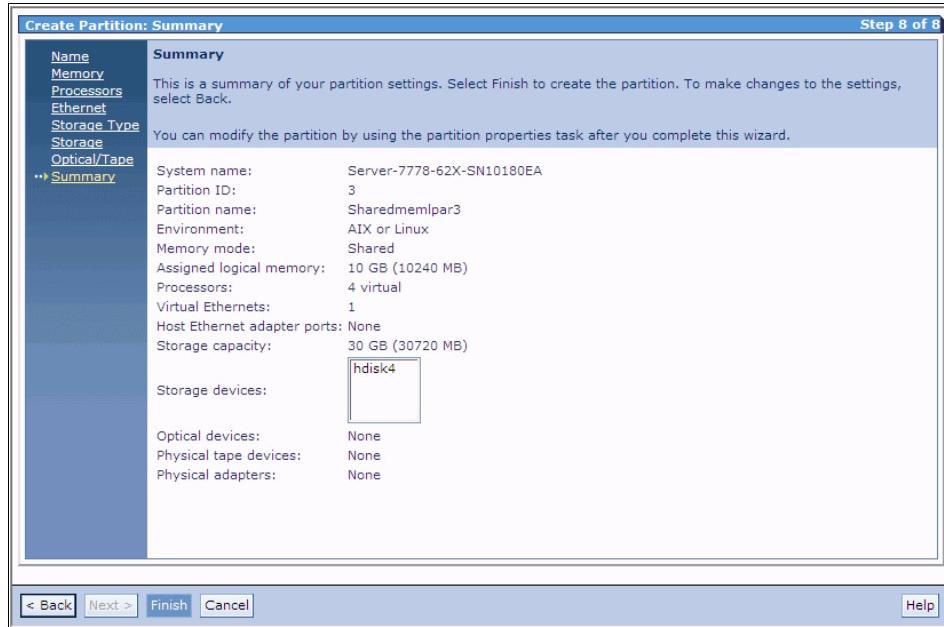


Figure 5-19 Summary of selections for a shared memory partition

Figure 5-20 on page 203 shows the View/Modify Partitions view with the new shared memory partition.

The screenshot shows the Integrated Virtualization Manager (IVM) interface. The left sidebar contains navigation links for Partition Management, I/O Adapter Management, Virtual Storage Management, IVM Management, System Plan Management, and Service Management. The main content area is titled "View/Modify Partitions" and displays system overview statistics and a table of existing partitions.

System Overview

Total system memory:	24 GB	Total processing units:	4
Memory available:	5.84 GB	Processing units available:	2.7
Reserved firmware memory:	928 MB	Processor pool utilization:	0.22 (5.4%)
Available shared memory pool size:	14.67 GB		
System attention LED:	Inactive		

Partition Details

Select	ID	Name	State	Uptime	Memory	Processors	Entitled Processing Units	Utilized Processing Units	Reference Code
<input type="checkbox"/>	1	is23-vios	Running	2.57 Days	2 GB	4	0.4	0.22	
<input type="checkbox"/>	2	IBMI_2	Not Activated		256 MB	1	0.5		00000000
<input type="checkbox"/>	3	Sharedmemlpar3	Not Activated		10 GB	4	0.4		00000000

Figure 5-20 View/Modify Partition window showing newly created shared memory partition

Figure 5-21 on page 204 shows the details of the shared memory pool indicating the new shared memory partition and the creation of lv00 in the paging storage pool supporting the partition Sharedmemlpar3.

Shared Memory Pool

A shared memory pool defines the amount of shared memory available on the system. Changes to the shared memory pool do not take effect until you click **Apply**.

Shared memory pool size: 15 GB (14.67 GB Available)
 Reserved firmware memory: 262 MB
 Total assigned logical memory: 10 GB
 Paging storage pool: AMS_Page_Pool (19.92 GB Available)

Property	Current	Pending
Assigned memory	15 GB	<input type="text" value="15"/> GB <input checked="" type="checkbox"/>
Maximum memory	15 GB	<input type="text" value="15"/> GB <input checked="" type="checkbox"/>

Delete Shared Memory Pool

Paging Space Devices - Advanced

A paging space device is a block storage device that is dedicated to the shared memory pool. When assigned to a shared memory partition, the paging space device provides paging space for the partition, as needed. When you create or modify a shared memory partition, IVM creates and manages the required paging space device for the partition automatically. However, you can define a specific paging space device for the shared memory pool, such as a physical volume. IVM can then assign the paging space device to a partition when you create it, if the device meets the appropriate requirements.

Click **Add** to define a new paging space device for the shared memory pool, or select a device and click **Remove**.

Add...	Remove	Select	Name ^	Storage Pool	Assigned Partition	Partition State	Size
<input type="radio"/>	lv00	AMS_Page_Pool	Sharedmem1par3 (3)	Not Activated	10 GB		

Apply **Reset**

Figure 5-21 shared memory pool with paging space assignments in paging pool

5.2.4 Shared memory partitions and dedicated paging devices

During the creation of the shared memory pool you have the option to create dedicated paging devices for shared memory partitions as detailed in 5.2.2, “Creating dedicated paging devices for partitions” on page 191.

These dedicated devices, if available, will be assigned by default if of adequate size to a shared memory partition when it is created. If the available dedicated devices have different sizes, the smallest size device that will meet the requirements of the assigned logical memory will be used. If the devices are all of equal size, the first available device will be assigned.

Figure 5-22 on page 205 shows the availability of two dedicated paging devices with sizes of 20GB and 30GB. These dedicated devices were added after the creation of the shared memory partition Sharedmem1par3.

Shared Memory Pool

A shared memory pool defines the amount of shared memory available on the system. Changes to the shared memory pool do not take effect until you click **Apply**.

Shared memory pool size: 15 GB (14.67 GB Available)
 Reserved firmware memory: 262 MB
 Total assigned logical memory: 10 GB
 Paging storage pool: AMS_Page_Pool (39.84 GB Available)

Property	Current	Pending
Assigned memory	15 GB	<input type="text" value="15"/> GB <input checked="" type="button" value="▼"/>
Maximum memory	15 GB	<input type="text" value="15"/> GB <input checked="" type="button" value="▼"/>

Delete Shared Memory Pool

Paging Space Devices - Advanced

A paging space device is a block storage device that is dedicated to the shared memory pool. When assigned to a shared memory partition, the paging space device provides paging space for the partition, as needed. When you create or modify a shared memory partition, IVM creates and manages the required paging space device for the partition automatically. However, you can define a specific paging space device for the shared memory pool, such as a physical volume. IVM can then assign the paging space device to a partition when you create it, if the device meets the appropriate requirements.

Click Add to define a new paging space device for the shared memory pool, or select a device and click Remove.

Add...	Remove	Select	Name ^	Storage Pool	Assigned Partition	Partition State	Size
<input type="button" value="Add..."/>	<input type="button" value="Remove"/>	<input type="radio"/>	lv00	AMS_Page_Pool	Sharedmem1par3 (3)	Not Activated	10 GB
		<input type="radio"/>	hdisk6			Not Available	20 GB
		<input type="radio"/>	hdisk2			Not Available	30 GB

Apply **Reset**

Figure 5-22 Shared memory pool view showing both types of paging devices

A new shared memory partition, Sharedmem1par4, was created with a logical memory value of 25GB. Figure 5-23 on page 206 shows this new partition and the assignment of hdisk2 as its dedicated paging device. Although the paging storage pool had over 39GB available, the default is to use dedicated paging devices when available. In this case the available hdisk2 with a size of 30GB was assigned to the partition Sharedmem1par4.

Integrated Virtualization Manager

Welcome padmin : baronlpar16.austin.ibm.com

Partition Management

- [View/Modify Partitions](#)
- [View/Modify System Properties](#)
- [View/Modify Shared Memory Pool](#)

I/O Adapter Management

- [View/Modify Host Ethernet Adapters](#)
- [View/Modify Virtual Ethernet](#)
- [View/Modify Physical Adapters](#)
- [View Virtual Fibre Channel](#)

Virtual Storage Management

- [View/Modify Virtual Storage](#)

IVM Management

- [View/Modify User Accounts](#)
- [View/Modify TCP/IP Settings](#)
- [Guided Setup](#)
- [Enter PowerVM Edition Key](#)

System Plan Management

- [Manage System Plans](#)

Service Management

- [Electronic Service Agent](#)
- [Service Focal Point](#)
 - [Manage Serviceable Events](#)
 - [Service Utilities](#)
 - [Create Serviceable Event](#)
 - [Manage Dumps](#)
 - [Collect VPD Information](#)
- [Updates](#)
- [Backup/Restore](#)
- [Application Logs](#)
- [Monitor Tasks](#)
- [Hardware Inventory](#)

Shared Memory Pool

A shared memory pool defines the amount of shared memory available on the system. Changes to the shared memory pool do not take effect until you click **Apply**.

Shared memory pool size: 15 GB (14.59 GB Available)
 Reserved firmware memory: 267 MB
 Total assigned logical memory: 35 GB
 Paging storage pool: AMS_Page_Pool (39.84 GB Available)

Property	Current	Pending
Assigned memory	15 GB	<input type="text" value="15"/> GB <input checked="" type="checkbox"/>
Maximum memory	15 GB	<input type="text" value="15"/> GB <input checked="" type="checkbox"/>

[Delete Shared Memory Pool](#)

Paging Space Devices - Advanced

A paging space device is a block storage device that is dedicated to the shared memory pool. When assigned to a shared memory partition, the paging space device provides paging space for the partition, as needed. When you create or modify a shared memory partition, IVM creates and manages the required paging space device for the partition automatically. However, you can define a specific paging space device for the shared memory pool, such as a physical volume. IVM can then assign the paging space device to a partition when you create it, if the device meets the appropriate requirements.

Click Add... to define a new paging space device for the shared memory pool, or select a device and click Remove.

Add...	Remove	Select	Name ^	Storage Pool	Assigned Partition	Partition State	Size
<input type="radio"/>		<input type="radio"/>	lv00	AMS_Page_Pool	Sharedmem1par3 (3)	Not Activated	10 GB
<input type="radio"/>		<input type="radio"/>	hdisk6			Not Available	20 GB
<input type="radio"/>		<input type="radio"/>	hdisk2		Sharedmem1par4 (4)	Not Activated	30 GB

[Apply](#) [Reset](#)

Figure 5-23 Shared memory pool view showing assigned dedicated paging device

Changing the maximum memory values of a shared memory partition can also cause a change from a paging pool logical volume to a dedicated paging device. Figure 5-24 on page 207 show the inactive partition Sharedmem1par3 that had its maximum memory value changed from 10GB to 15GB. When this change was made the paging space changed from 10GB lv00 in the pool AMS_Page_Pool to the 30GB hdisk6. Also note the informational message indicating current and pending values are not synchronized. A partition activation will complete the process.

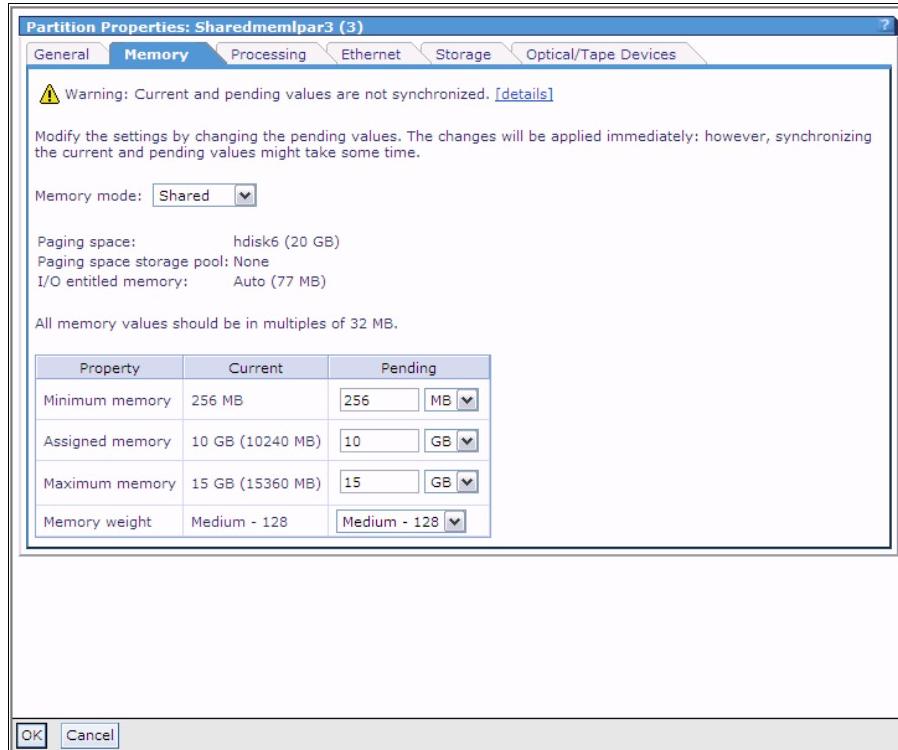


Figure 5-24 Partition memory properties showing maximum memory and paging space changes

Figure 5-25 on page 208 shows the shared memory pool indicating the changes to the paging device used for the partition Sharedmem1par3 when the maximum memory values were changed.

Integrated Virtualization Manager

Welcome padmin : baronlpar16.austin.ibm.com

Partition Management

- [View/Modify Partitions](#)
- [View/Modify System Properties](#)
- [View/Modify Shared Memory Pool](#)

I/O Adapter Management

- [View/Modify Host Ethernet Adapters](#)
- [View/Modify Virtual Ethernet](#)
- [View/Modify Physical Adapters](#)
- [View Virtual Fibre Channel](#)

Virtual Storage Management

- [View/Modify Virtual Storage](#)

IVM Management

- [View/Modify User Accounts](#)
- [View/Modify TCP/IP Settings](#)
- [Guided Setup](#)
- [Enter PowerVM Edition Key](#)

System Plan Management

- [Manage System Plans](#)

Service Management

- [Electronic Service Agent](#)
- [Service Focal Point](#)
 - [Manage Serviceable Events](#)
 - [Service Utilities](#)
 - [Create Serviceable Event](#)
 - [Manage Dumps](#)
 - [Collect VPD Information](#)
- [Updates](#)
- [Backup/Restore](#)
- [Application Logs](#)
- [Monitor Tasks](#)
- [Hardware Inventory](#)

Shared Memory Pool

A shared memory pool defines the amount of shared memory available on the system. Changes to the shared memory pool do not take effect until you click **Apply**.

Shared memory pool size: 15 GB (14.66 GB Available)
 Reserved firmware memory: 267 MB
 Total assigned logical memory: 35 GB
 Paging storage pool: AMS_Page_Pool (49.84 GB Available)

Property	Current	Pending
Assigned memory	15 GB	<input type="text" value="15"/> GB
Maximum memory	15 GB	<input type="text" value="15"/> GB

[Delete Shared Memory Pool](#)

Paging Space Devices - Advanced

A paging space device is a block storage device that is dedicated to the shared memory pool. When assigned to a shared memory partition, the paging space device provides paging space for the partition, as needed. When you create or modify a shared memory partition, IVM creates and manages the required paging space device for the partition automatically. However, you can define a specific paging space device for the shared memory pool, such as a physical volume. IVM can then assign the paging space device to a partition when you create it, if the device meets the appropriate requirements.

Click Add to define a new paging space device for the shared memory pool, or select a device and click Remove.

Add...	Remove				
Select	Name ^	Storage Pool	Assigned Partition	Partition State	Size
<input type="radio"/>	hdisk6		Sharedmem1par3 (3)	Not Activated	20 GB
<input type="radio"/>	hdisk2		Sharedmem1par4 (4)	Not Activated	30 GB

[Apply](#) [Reset](#)

Figure 5-25 shared memory pool after partition maximum memory values changed

5.2.5 Active Memory Sharing DLPAR operations

Dynamic logical partition (DLPAR) operations can be performed on both the shared memory pool and shared memory partition logical memory assignments. The assigned memory in shared memory pool can be DLPARed up to its maximum value, and the memory pool maximum value can be dynamically increase up to the available limits of the physical memory minus firmware requirements. The shared memory partition can be altered between the minimum and maximum values, as shown in the partition properties under the memory tab.

5.3 Active Memory Sharing summary

Active Memory Sharing provides the ability to better utilize the memory *and* CPU resources available on a IBM BladeCenter JS23 or JS43. However the successful implementation requires a complete understanding of current or planned workloads and the proper matching of those workload in the right combinations. Improper matching will result in contention for memory resources and excessive paging by the VIO Server in an attempt to service the partitions memory needs.

Dedicated paging devices are the recommended method for providing paging space for the shared memory pool. If a mix of dedicated and logical volumes are used additional planning is required to determine when to add the dedicated devices in the sequence of creating shared memory partitions or changing existing dedicated memory partitions.



IBM AIX V6.1 installation

IBM AIX can be installed native on IBM BladeCenter JS23 and JS43 Express or in a client partition of IBM PowerVM. This chapter describes in details the installation on a logical partition and has the following sections:

- ▶ “Create a virtual media library” on page 212
- ▶ “Prepare the PowerVM client partition” on page 218
- ▶ “Install AIX 6.1 in a logical partition of the Virtual IO Server” on page 231

6.1 Install IBM AIX 6.1 in a PowerVM client partition

This section assumes that you have already installed VIOS 1.5.2.1 or any later version (latest version is V2.1.1) on the blade and performed the initial configuration. In case this was not done, go to 4.2, “VIOS system management using IVM” on page 83.

To install IBM AIX 6.1 in a client partition it is necessary to first create the client partition with the IVM before you can start with the installation of AIX. This chapter is divided into the following parts:

- ▶ “Create a virtual media library” on page 212
- ▶ “Prepare the PowerVM client partition” on page 218
- ▶ “Install AIX 6.1 in a logical partition of the Virtual IO Server” on page 231

6.1.1 Create a virtual media library

A virtual media library is a new feature that was introduced in the Virtual IO Server version 1.5. It allows to store images from CDs and DVDs in a logical volume. These images can be mounted to virtual optical devices that are assigned to partitions.

A media library is created for the AIX install DVD that is used to install the first AIX partition. This section describes how to create a storage pool that will be used to create the logical volume for the media library. Afterwards we will describe how to add a DVD image, from the AIX DVD, to the media library created.

Perform the following steps to set up a media library.

1. Click **View/Modify Virtual Storage** in the left menu under Virtual Storage Management. Then click the **Storage Pools** register card. Create a new storage pool by clicking **Create Storage Pool**. See Figure 6-1.

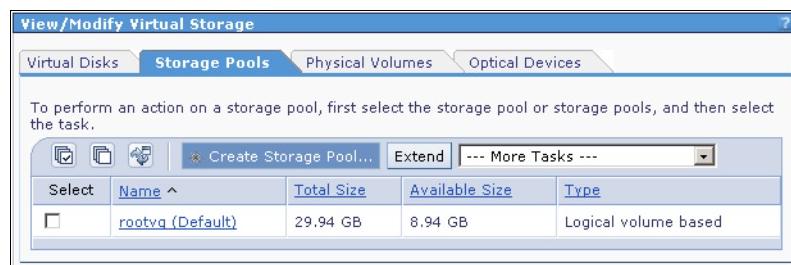


Figure 6-1 Media library - create storage pool

- Specify the name of the storage pool and select the physical volumes that will be assigned to this storage pool. Figure 6-2 shows that we used the name STG-Pool-Media1. The type of the volume group is *logical volume base*. This allows to increase the space of the media library when needed. Physical volume hdsik3 is assigned to this pool. Click **OK**.

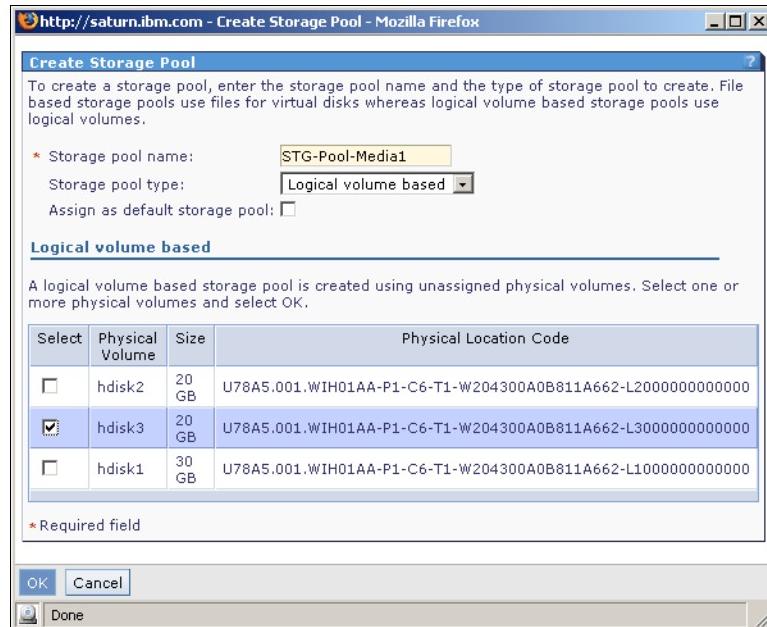


Figure 6-2 Media library - select the physical volume

- The storage pool was created. Now select the **Optical Devices** register card. See Figure 6-3.

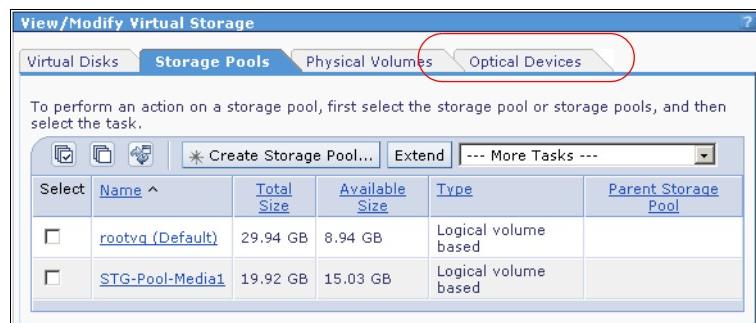


Figure 6-3 Media library - new storage pool

4. Click **Create Library**. See Figure 6-4.



Figure 6-4 Media library - create library

5. Specify the storage pool that will contain the logical volume with the media library and the size of the media library. We used the volume group created in step 1 on page 212. The initial size was set to hold the AIX 6.1 DVD with a size of approximately 3.6 GB. See Figure 6-5. Click **OK**.

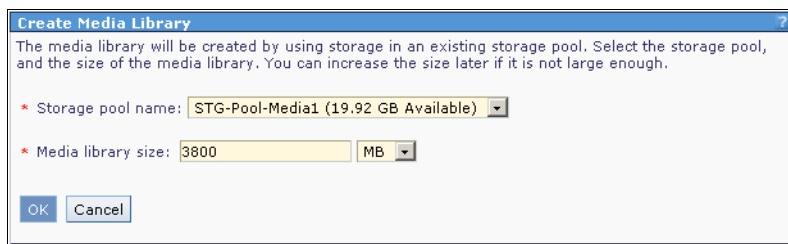


Figure 6-5 Media library - specify library size and volume group

- It takes a moment to create the library volume and file. After that is done, return to the panel shown in Figure 6-6. Click **Add Media** to create an image from the AIX DVD.

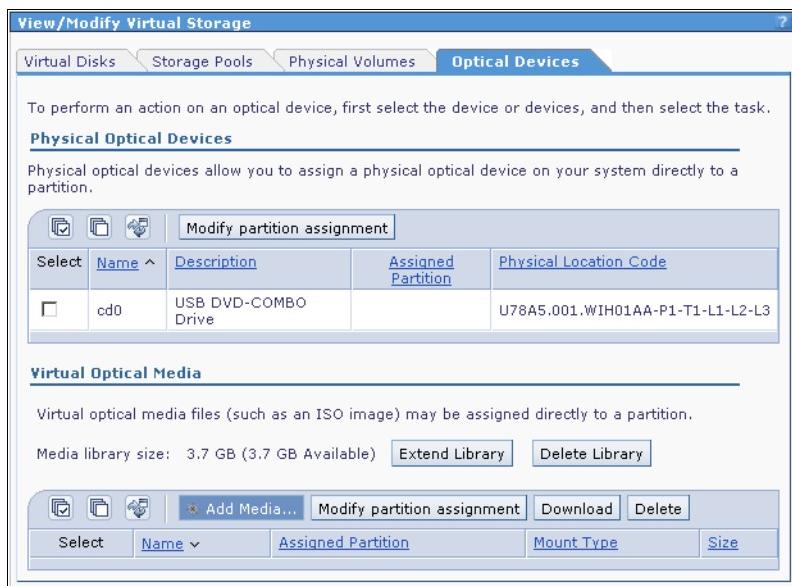


Figure 6-6 Media library - add media

- The add media dialog starts and guides you through the process of adding medias to the library. Click **Import from physical optical device** to get the list of available physical optical devices that you can use to import the media. Specify the media type of the new media. Choose between read-only and read/write mode of the new media. Now enter the name of the new media. This will be used as file name to store the content from the CD or DVD. Click **OK** to start the copy process.

Note: Do not use a media name that contains spaces in the name. This will lead to an error message like the one shown in Figure 6-7.

Problems occurred while processing the data. A summary of all problems for this page are listed below. Additional details for each problem may be located next to the field causing the problem.



Too many parameters.

Usage: mkvopt -name FileName {-size FileSize | -dev SourceDevice
-file SourceFile} [-ro]

Figure 6-7 Error message using a wrong media name

You may look at existing media files in /var/vio/VMLibrary. The last step on this page is the specification of the optical device that contains the CD or DVD to copy into the library. Figure 6-8 shows the optical device that is located in the media tray of the IBM BladeCenter H chassis. The remote media optical device uses the location code U78A5.001.WIH01AA-P1-T1-L1-L1.

We used the internal optical device of the BladeCenter chassis to copy the data from the IBM AIX 6.1 DVD. This took approximately two hours.

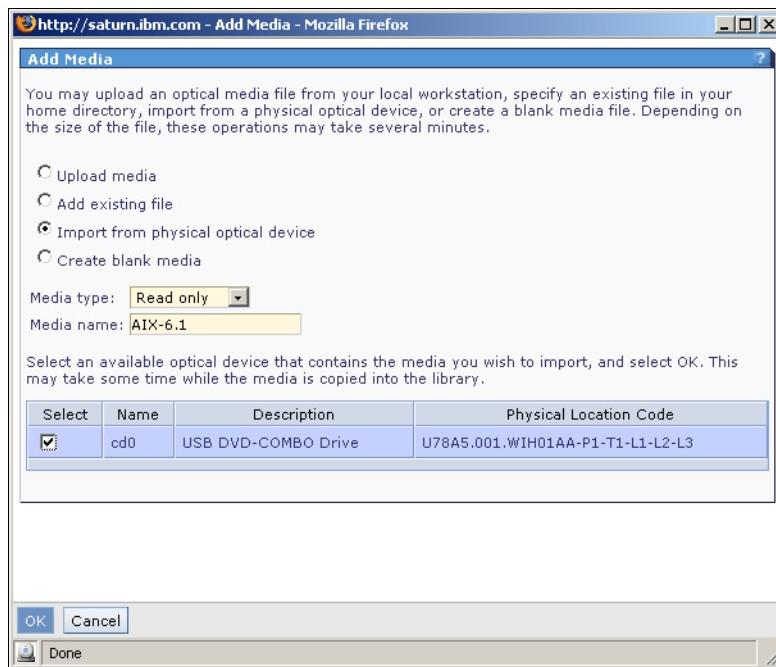


Figure 6-8 Media library - specify details about media creation

8. The copy task will take some time. You may close the dialog browser window and proceed with other tasks in the meantime. At any time you can check whether the task has completed or failed by using the Monitor Task function.

This function can be reached with **Monitor Task** before you close your browser window or from the main window's left-hand navigation under **Service Management → Monitor Task**. See Figure 6-9.

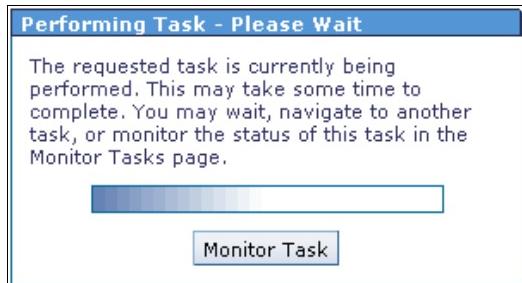


Figure 6-9 Media library - performing task

9. After closing the browser window of the add media dialog, you return to the view shown in Figure 6-10. The new media is already listed here. Clicking **Refresh** updates the size information during the copy operation.

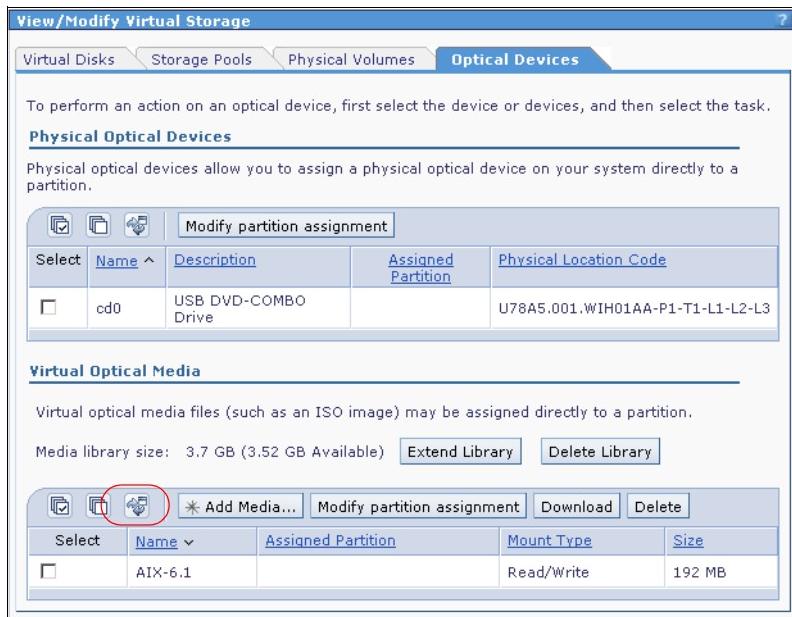


Figure 6-10 Media library - newly created media

The media library is now ready to be used. When the free space in the library is not sufficient for new CD images, expand the media library. This can be done at any time.

6.1.2 Prepare the PowerVM client partition

Perform the following steps to create a client partition with the Integrated Virtualization Manager (IVM) of the Virtual I/O Server.

1. Use your preferred Web browser and enter the host name or IP address of the IVM. That is the address configured in 4.3.3, “Initial network setup” on page 87.

A Web page comes up that allows you to log in. Use the default account that was created during setup when you had not yet created your own account.

The default account is called *padmin* and uses the password *padmin*.

Figure 6-11 shows the logon window of the IVM. Click **Login**.

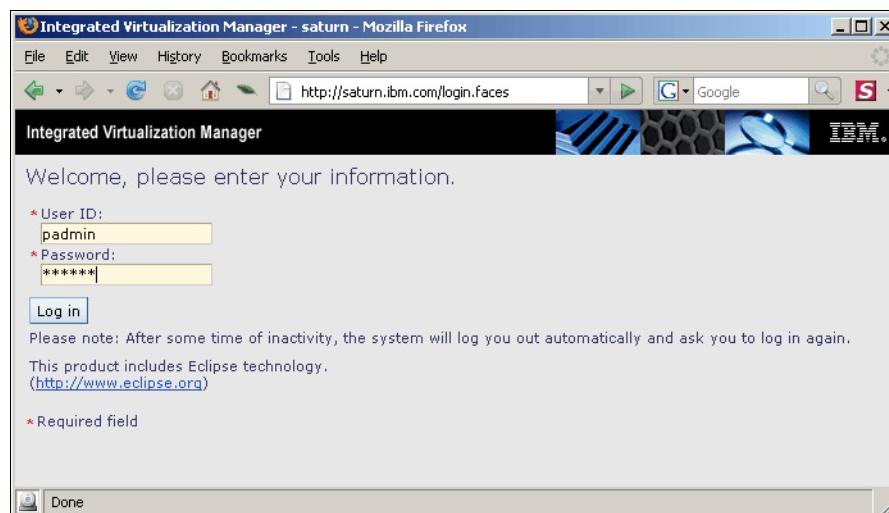


Figure 6-11 IVM login window

2. Depending on the setup of your IVM, you will be at the Guided Setup or on the View/Modify Partitions page. Figure 6-12 shows the usual page that you see after logon when the IVM is fully configured.

The screenshot shows a Mozilla Firefox browser window titled "Integrated Virtualization Manager - saturn - Mozilla Firefox". The URL is "http://saturn.ibm.com/main.faces". The main content area is titled "View/Modify Partitions" and contains the following information:

To perform an action on a partition, first select the partition or partitions, and then select the task.

System Overview

Total system memory:	4 GB	Total processing units:	4
Memory available:	2.7 GB	Processing units available:	3.6
Reserved firmware memory:	304 MB	Processor pool utilization:	0.00 (0.1%)
System attention LED:	Inactive		

Partition Details

Select	ID	Name	State	Uptime	Memory	Processors	Entitled Processing Units	Utilized Processing Units	Reference Code
<input type="checkbox"/>	1	10-0059A	Running	14.18 Hours	1 GB	4	0.4	0.00	

Figure 6-12 View/Modify Partitions page after logon

3. Verify that you have your storage available to the VIOS. Click **View/Modify Virtual Storage** in the left menu under Virtual Storage Management. See Figure 6-13.

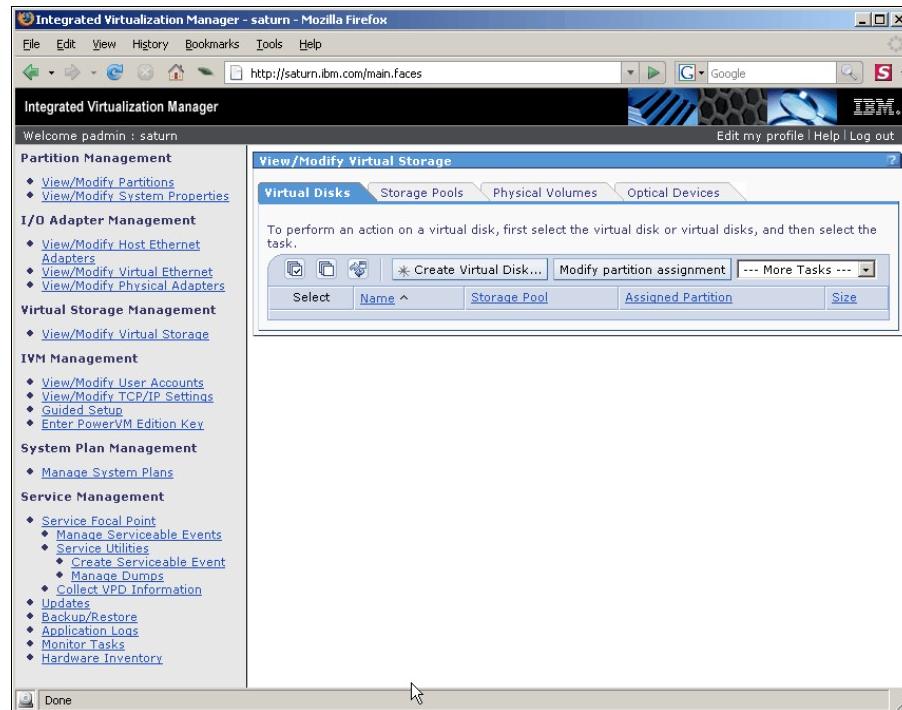


Figure 6-13 View/Modify Storage

- Click the **View/Modify Storage** page on the Physical Volumes tab to see a list of available hard drives to the VIOS. Verify that the expected drives are available. See Figure 6-14.

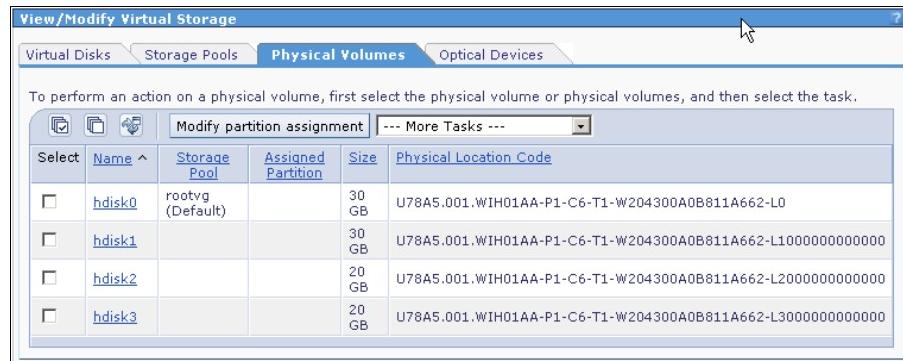


Figure 6-14 Available physical volumes

Figure 6-15 shows that there are four physical volumes available. They are all located on a DS4800. HDISK0 and HDISK1 are used for the VIOS itself. HDISK2 will be used for AIX client partitions that will be created in the next steps.

Click **View/Modify Partitions** under Partition Management. Then click **Create Partition** as shown in Figure 6-15.

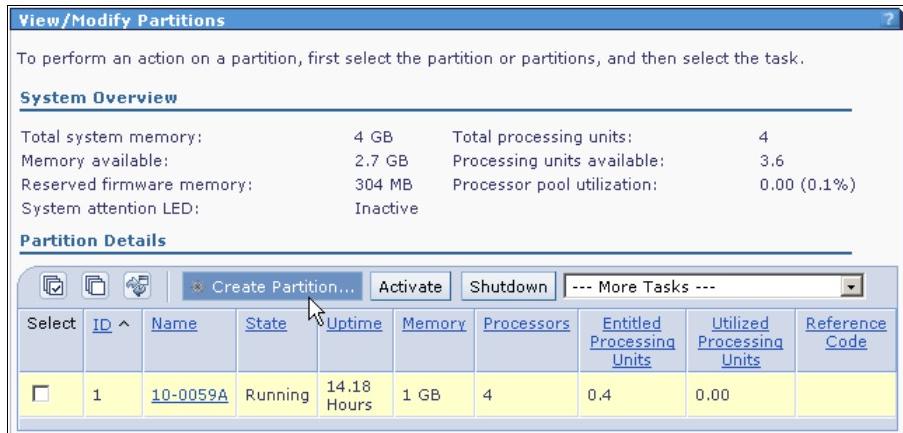


Figure 6-15 View/ Modify Partition - Create Partition

A dialog will open that guides you through the process of partition creation.

5. Specify the name and the type of the partition. The name is used to identify the partition, especially when partition mobility is later used. Using a host name might be an option here. In Figure 6-16 we chose the host name as partition name. The type can be either AIX/Linux or i5/OS. Choose the type according to the OS you plan to install. We chose AIX/Linux for this AIX partition. Click **Next** to proceed.

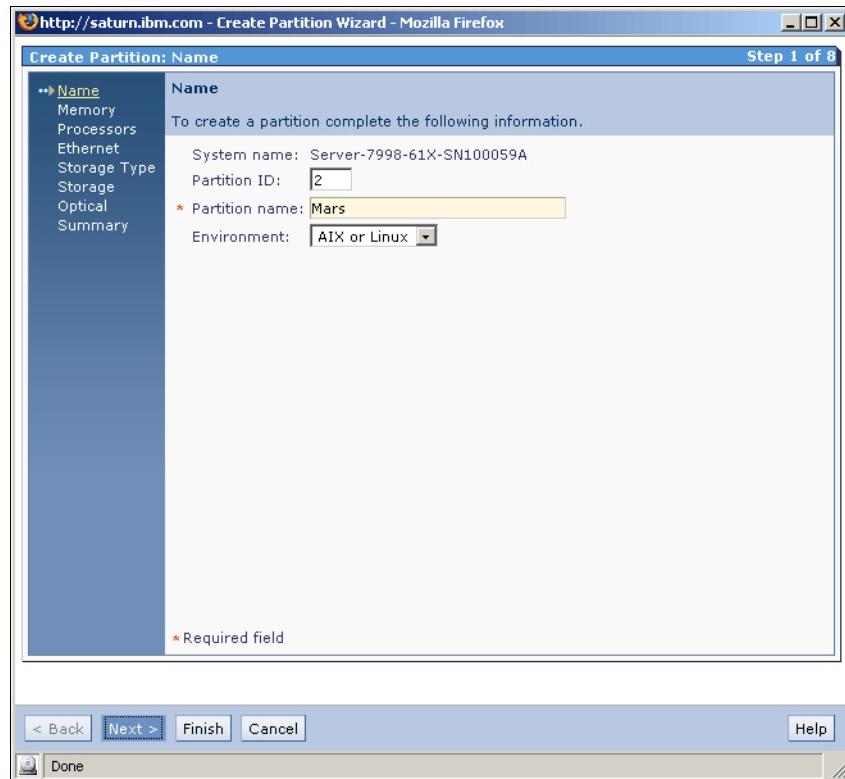


Figure 6-16 Create partition - define name

6. Define the amount of memory that will be assigned to the partition. In Figure 6-17 we chose 1 GB. Click **Next** to proceed.

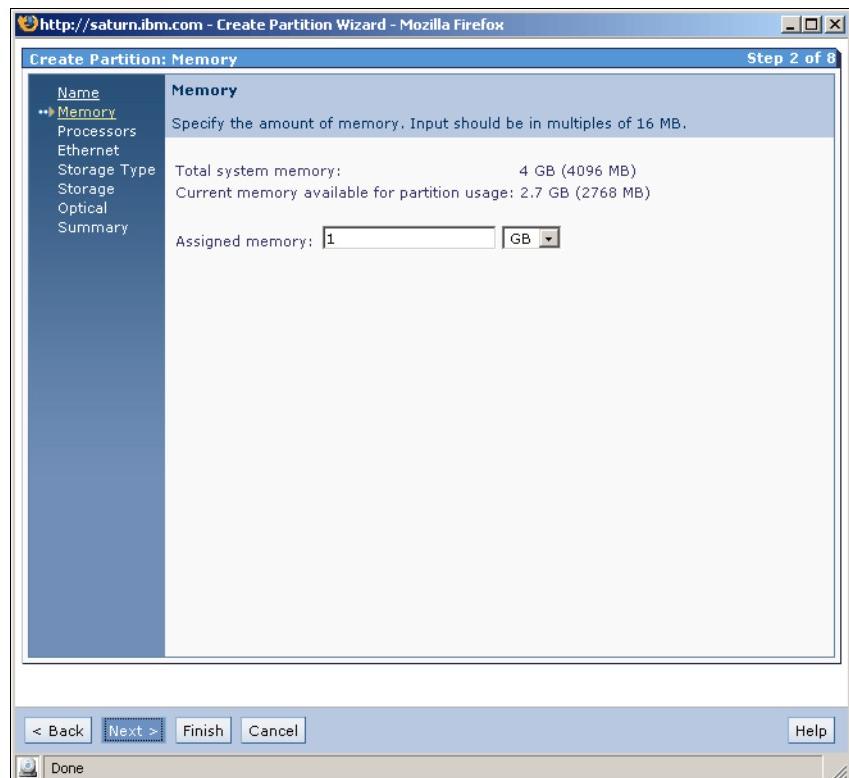


Figure 6-17 Create partition - configure the amount of memory

7. Choose the number of CPUs that will be used by the partition. You have to decide whether to use dedicated or shared CPUs. When a dedicated CPU is used, no load can be moved to other currently free CPUs because this may lead to a performance issue. In Figure 6-18 you see that we configured two CPUs and shared processor mode. Click **Next** to proceed.

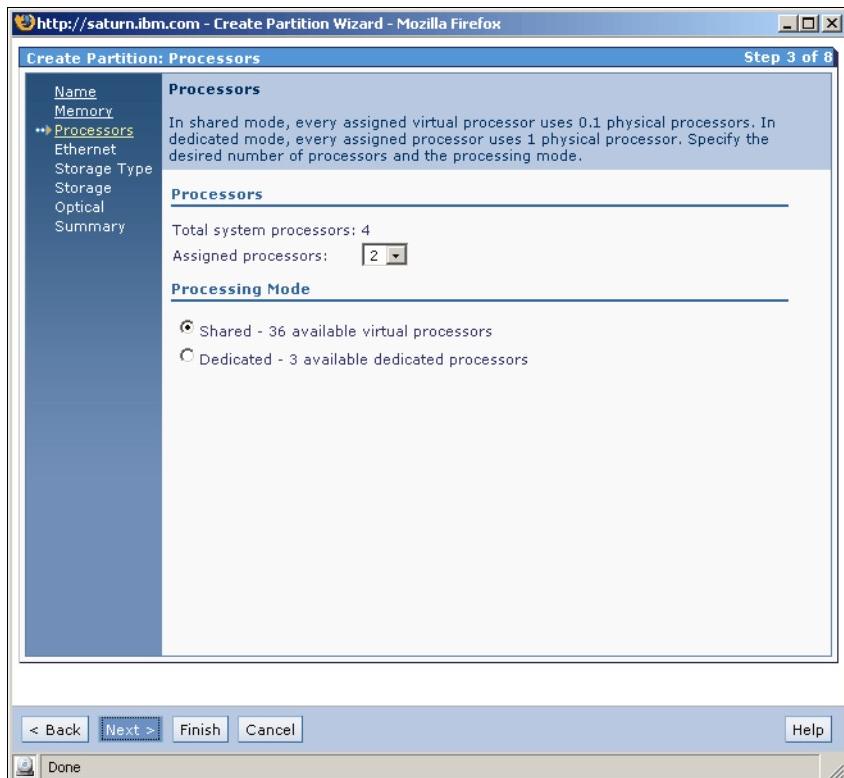


Figure 6-18 Create partition - CPU configuration

8. Depending on the setup of the network in the BladeCenter chassis and in the Virtual IO Server, you may have different settings. When you define host Ethernet adapters to the partition then no Ethernet bridge in the Virtual IO Server is required. The disadvantage is that you are not able to use partition mobility. For partition mobility it is a requirement that the partition be configured with virtual Ethernet adapters. The Ethernet bridge is configured in the Virtual IO Server. For more details about this see 4.5.2, “Virtual Ethernet Adapters and SEA” on page 103.

As shown in Figure 6-19, we chose one virtual Ethernet adapter. Click **Next** to proceed.

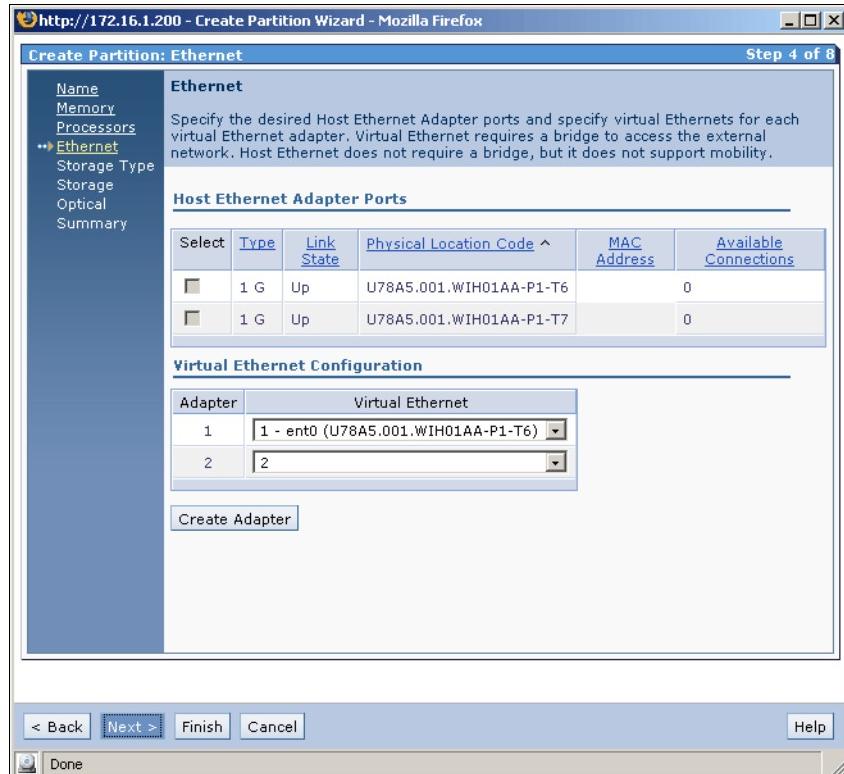


Figure 6-19 Create partition - ethernet network

9. Set up the storage type you plan to use. There are three different options available. You may use volume group or file-based storage. In addition there is an option to use a dedicated physical volume for the partition. Depending on the type of storage subsystem, there are limitations about the maximum number of physical volumes you may have per host adapter or host adapter group. In case of the DS4800 you can have up to 32 logical drives from the DS4000 assigned to a host or host group. In our scenario we used a dedicated physical volume for each partition.

As shown in Figure 6-20, select **Assign existing virtual disks and physical volumes**. Click **Next** to proceed.

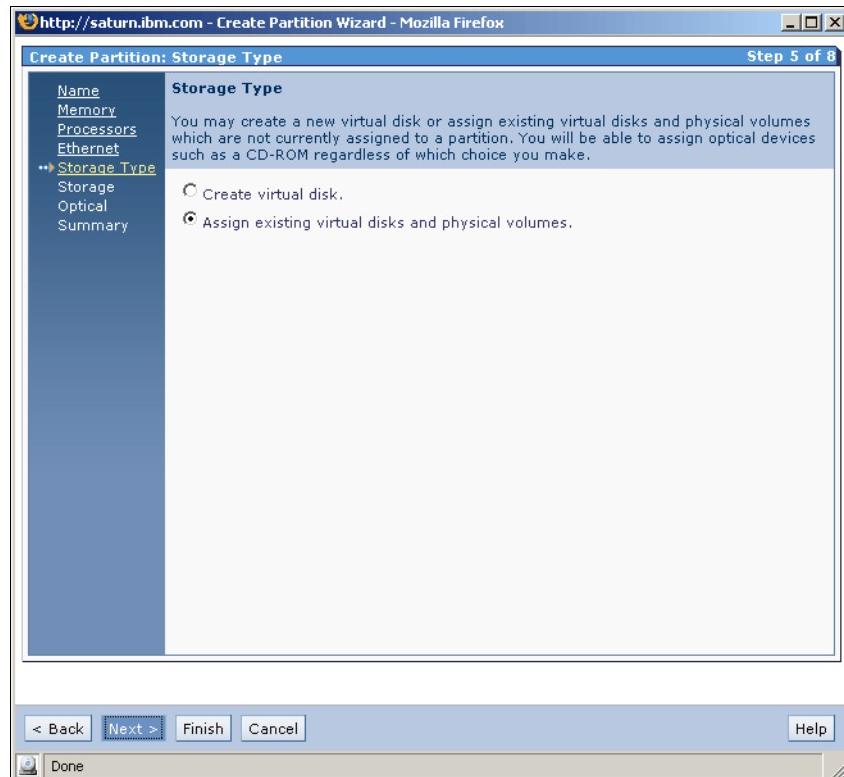


Figure 6-20 Create partition - storage type

10. Select the physical volume or volumes that need to be available to the partition. Figure 6-21 shows the section of hdisk1. Click **Next** to proceed.

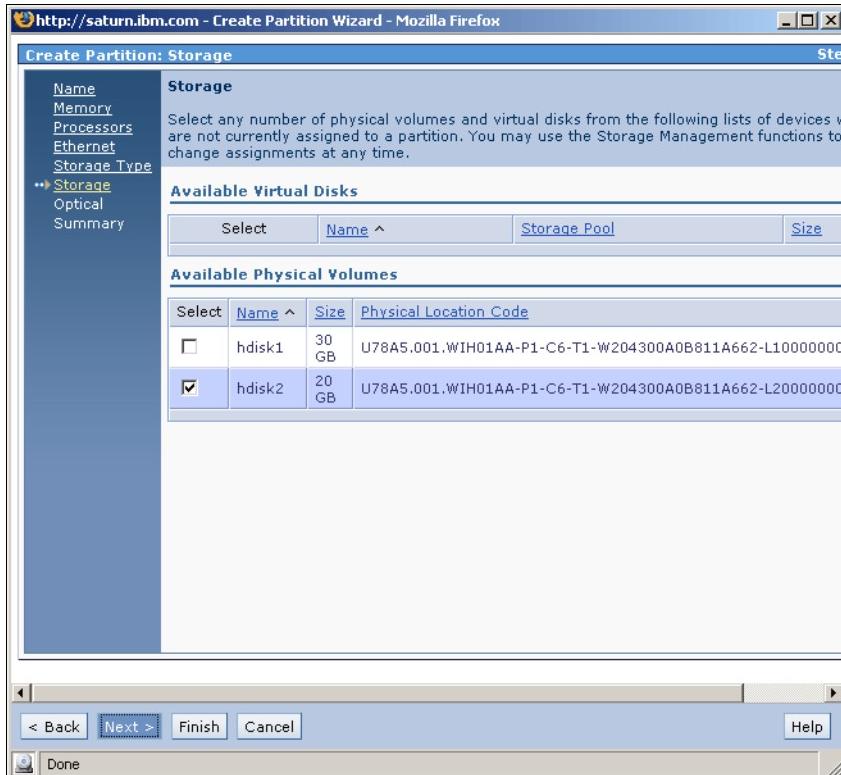


Figure 6-21 Create partition - select physical volumes

11. In the optical section of the partition creation process you can define the CD-ROM drives that will be used by the partition. Two options are possible:

- Physical drive attached to the partition
- Virtual drive attached to the partition

There might be multiple physical CDROM drives available. Use the location code to differentiate between the CDROM drives:

- U78A5.001.WIH01AA-P1-T1-L1-L2-L3 - CDROM drive in the media tray
- U78A5.001.WIH01AA-P1-T1-L1-L1 - Remote media CDROM drive

The local drive installed in the media tray of the IBM BladeCenter chassis is identified by the location code U78A5.001.WIH01AA-P1-T1-L1-L2-L3. The CDROM drive that has the location code U78A5.001.WIH01AA-P1-T1-L1-L1

is the CDROM drive that is provided via the Remote Control Web interface of the Advanced Management Module.

Note: When you attach the media tray of the BladeCenter chassis to a blade that is already up and running you may have to issue **cfgdev** on the command line of the Virtual IO Server to get it recognized by VIOS.

Virtual CDROM drives are used to mount CDs that are placed in the media library. See 4.6.2, “Storage pools” on page 124 and 4.6.4, “Optical and Tape devices” on page 132.

The current setup uses an AIX CD that was placed in the media library. The first virtual optical device is checked by default. Click the **Modify** link to select an image from the media library. See Figure 6-22.

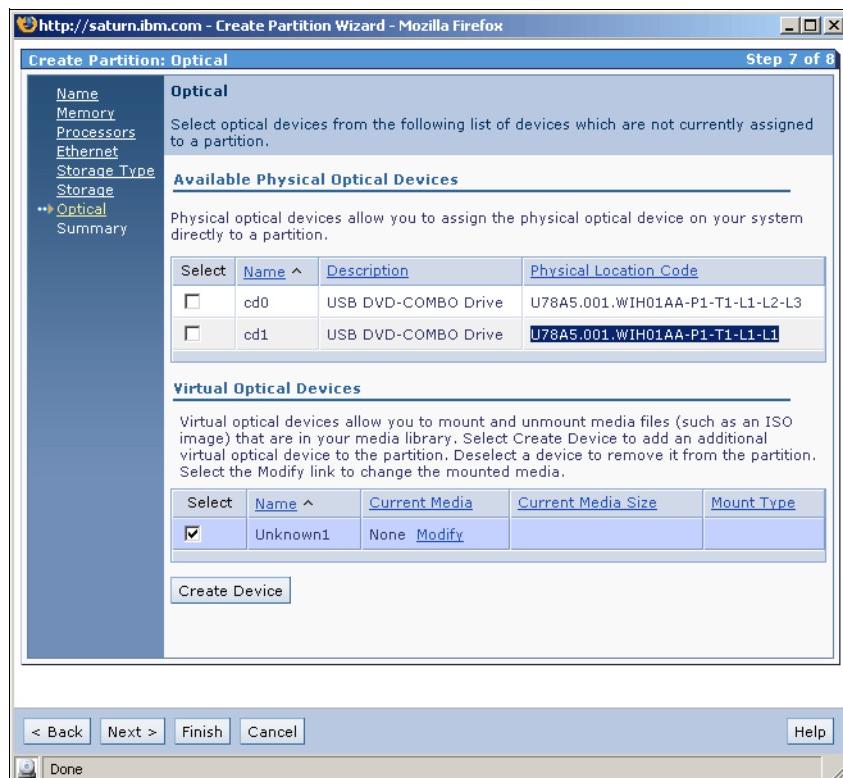


Figure 6-22 Create partition - modify virtual optical device

12. Change the selected media from *none* to AIX-6.1 and click **OK**. See Figure 6-23.



Figure 6-23 Create partition - modify current media of virtual optical device

13. Click **Next** to see an overview of the setting of the new partition. See Figure 6-24.

A screenshot of a web-based partition creation interface. The title bar says "http://saturn.ibm.com - Create Partition Wizard - Mozilla Firefox" and "Step 7 of 8". On the left is a sidebar with links: Name, Memory, Processors, Ethernet, Storage Type, Storage, and "Optical" (which is selected and highlighted in blue).
Available Physical Optical Devices:
Physical optical devices allow you to assign the physical optical device on your system directly to a partition.

Select	Name ^	Description	Physical Location Code
<input type="checkbox"/>	cd0	USB DVD-COMBO Drive	U78A5.001.WIH01AA-P1-T1-L1-L2-L3
<input type="checkbox"/>	cd1	USB DVD-COMBO Drive	U78A5.001.WIH01AA-P1-T1-L1-L1

Virtual Optical Devices:
Virtual optical devices allow you to mount and unmount media files (such as an ISO image) that are in your media library. Select Create Device to add an additional virtual optical device to the partition. Deselect a device to remove it from the partition. Select the Modify link to change the mounted media.

Select	Name ^	Current Media	Current Media Size	Mount Type
<input checked="" type="checkbox"/>	Unknown1	AIX-6.1 Modify	3.59 GB	Read only

At the bottom are navigation buttons: < Back, Next >, Finish, Cancel, Help, and Done.

Figure 6-24 Create partition - virtual optical device

14. Verify your setting and click **Finish** to create a partition with the settings you defined. See Figure 6-25.

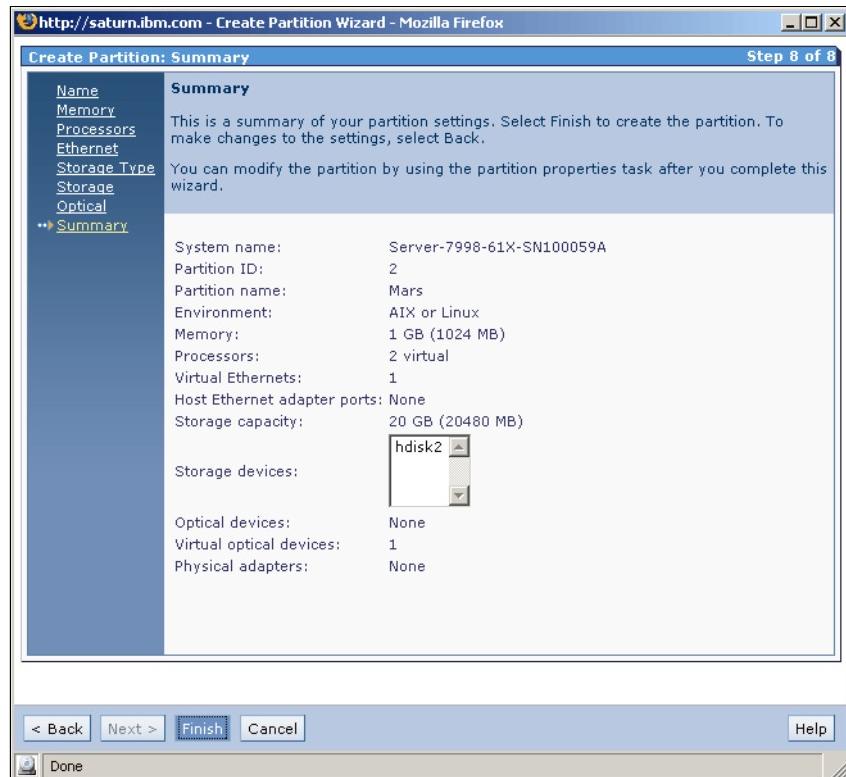


Figure 6-25 Create partition - summary

15. The new partition will be listed under View/Modify Partitions, as shown in Figure 6-26.

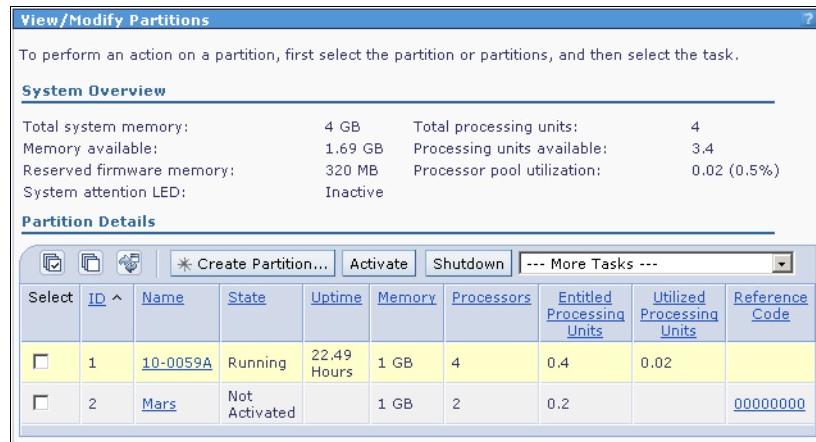


Figure 6-26 Newly created AIX/Linux partition

The preparation of the partition is done. Proceed with the installation of AIX in the newly created partition.

6.1.3 Install AIX 6.1 in a logical partition of the Virtual IO Server

The previous sections described how to prepare the media library that contains the AIX 6.1 DVD image that will be used to install the first logical partition, and how to create a logical partition. This section describes the activation of the logical partition and the installation of AIX 6.1 from a virtual optical device. Follow the outlined steps:

1. To activate the partition, click the check box of the partition and click **Activate**. See Figure 6-27.

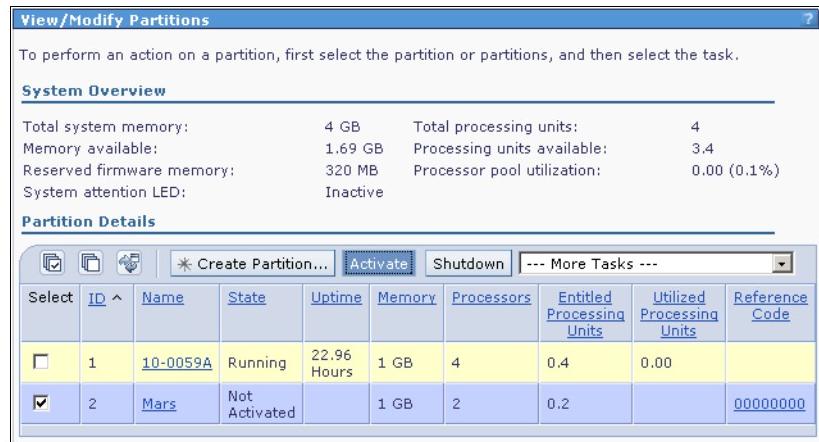


Figure 6-27 Activate a partition

2. Confirm the activation of the partition with **OK** as shown in Figure 6-28.

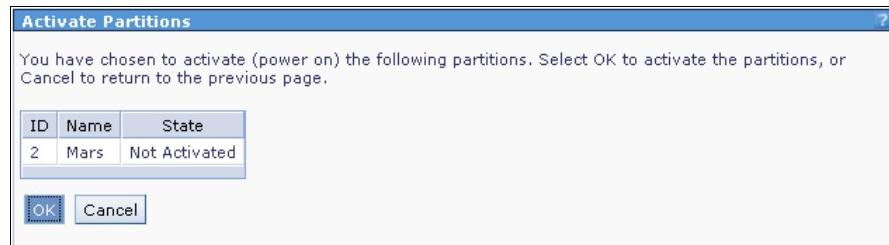


Figure 6-28 Confirm partition activation

3. The status of the partition has changed to *running*. Select **Open Terminal** from the More Tasks drop-down list box to open a terminal connected to the selected partition. See Figure 6-29.

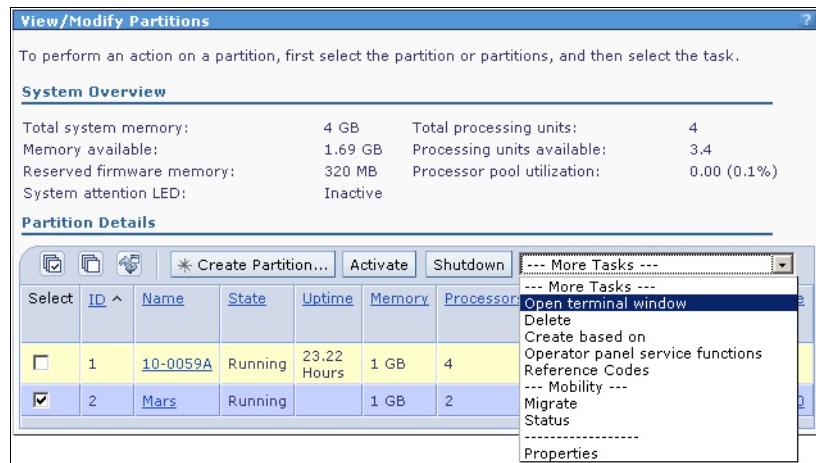


Figure 6-29 Open a virtual terminal to the partition

4. Authenticate on the Virtual IO Server to get the virtual terminal connected. You may use the account padmin with the default password padmin here in case you have not yet created your own account. After the authentication is done, a message will be shown that the terminal has connected, as shown in Figure 6-30.

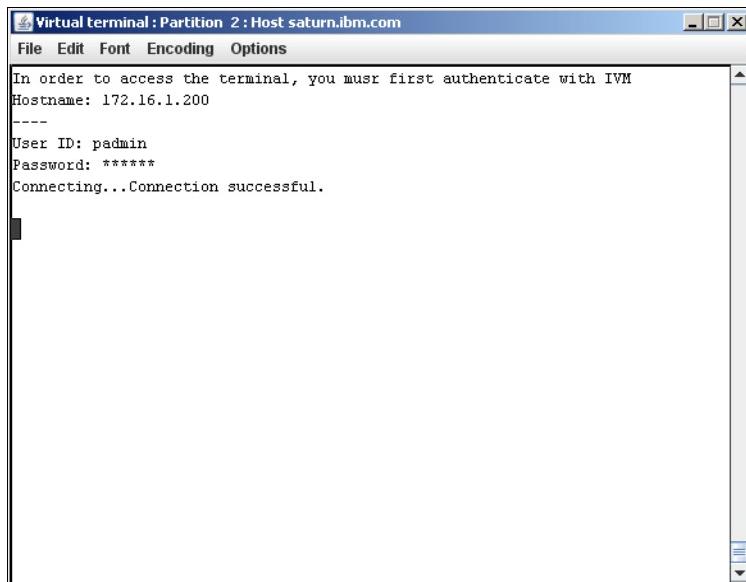


Figure 6-30 Virtual terminal connection

On the virtual terminal you will see the POST of the partition with the possibility to enter the SMS menu. There is no change required in this stage. The partition will boot from the assigned virtual optical device.

5. When the partition has booted from the virtual optical device, you see a few messages, as shown in Example 6-1.

Example 6-1 AIX install kernel load

```
Elapsed time since release of system processors: 1665 mins 26 secs
```

```
Welcome to AIX.  
boot image timestamp: 23:19 10/31  
The current time and date: 01:23:30 06/11/2008  
processor count: 2; memory size: 1024MB; kernel size: 26145029  
boot device:  
/vdevice/v-scsi@30000002/disk@8200000000000000:\ppc\chrp\bootfile.exe  
kernel debugger setting: enabled
```

```
AIX Version 6.1
Starting NODE#000 physical CPU#001 as logical CPU#001... done.
Starting NODE#000 physical CPU#002 as logical CPU#002... done.
Starting NODE#000 physical CPU#003 as logical CPU#003... done.
Preserving 126407 bytes of symbol table [/usr/lib/drivers/hd_pin]
Preserving 199549 bytes of symbol table [/usr/lib/drivers/hd_pin_bot]
```

6. Define the current virtual terminal as system console by entering **1**. Click Enter to proceed; see Example 6-2. Depending on the console you are using, you may need to also enter **F1** or **2**.

Example 6-2 Select the system console

```
Preserving 199549 bytes of symbol table [/usr/lib/drivers/hd_pin_bot]
```

***** Please define the System Console. *****

Type a 1 and press Enter to use this terminal as the system console.

Pour definir ce terminal comme console système, appuyez sur 1 puis sur Entrée.

Taste 1 und anschliessend die Eingabetaste druecken, um diese Datenstation als Systemkonsole zu verwenden.

Premere il tasto 1 ed Invio per usare questo terminal come console.

Escriba 1 y pulse Intro para utilizar esta terminal como consola del sistema.

Escriviu 1 i premeu Intro per utilitzar aquest terminal com a consola del sistema.

Digite um 1 e pressione Enter para utilizar este terminal como console do sistema.

7. Select the number of the language that you would like to use during the installation of IBM AIX 6.1. You can define the language of the operating system that now gets installed in a later step. Click Enter to proceed; see Example 6-3.

Example 6-3 Select the language used during installation

```
>>> 1 Type 1 and press Enter to have English during install.
      2 Entreu 2 i premeu Intro per veure la instal·lació en català.
      3 Entrez 3 pour effectuer l'installation en français.
      4 Für Installation in deutscher Sprache 4 eingeben
          und die Eingabetaste drücken.
      5 Immettere 5 e premere Invio per l'installazione in Italiano.
      6 Digite 6 e pressione Enter para usar Português na instalação.
      7 Escriba 7 y pulse Intro para la instalación en español.
```

```
88 Help ?
```

```
>>> Choice [1]:
```

8. Modify required settings such as language or time zone and proceed with the installation by entering **1** followed by Enter, as shown in Example 6-4.

Example 6-4 AIX installation summary

Overwrite Installation Summary

```
Disks: hdisk0
Cultural Convention: en_US
Language: en_US
Keyboard: en_US
JFS2 File Systems Created: Yes
Graphics Software: Yes
System Management Client Software: Yes
Enable System Backups to install any system: Yes
```

Optional Software being installed:

```
>>> 1 Continue with Install
```

```
+-----+
88 Help ? | WARNING: Base Operating System Installation will
99 Previous Menu | destroy or impair recovery of ALL data on the
| destination disk hdisk0.
```

```
>>> Choice [1]:
```

9. The copy process starts after you click **1** followed by the Enter key.
10. After the installation is done, a reboot of the partition is performed. Then select your terminal type, as shown in Example 6-5.

Example 6-5 Select the terminal type you are using

Set Terminal Type

The terminal is not properly initialized. Please enter a terminal type and press Enter. Some terminal types are not supported in non-English languages.

ibm3101	tvi912	vt330	aixterm
ibm3151	tvi920	vt340	dtterm
ibm3161	tvi925	wyse30	xterm
ibm3162	tvi950	wyse50	lft
ibm3163	vs100	wyse60	sun

ibm3164	vt100	wyse100
ibmpc	vt320	wyse350

+-----Messages-----

If the next screen is unreadable, press Break (Ctrl-c) to return to this screen.

>>> Choice []:

11. Select **Show Installed License Agreements** and click Enter to read the license agreement; see Example 6-6.

Example 6-6 License agreement menu

Software License Agreements

Move cursor to desired item and press Enter.

Show Installed License Agreements
Accept License Agreements

F1=Help	F2=Refresh	F3=Cancel	Esc+8=Image
Esc+9=Shell	Esc+0=Exit	Enter=Do	

12. Select the software package from which you would like to read the license agreements. The default is to show all license agreements. Click Enter to start showing the license text; see Example 6-7.

Example 6-7 Select a software license

Show Installed License Agreements

Type or select values in entry fields.
Press Enter AFTER making all desired changes.

[Entry Fields]	
* SOFTWARE name	[all] +
SHOW license agreement text?	yes +

F1=Help	F2=Refresh	F3=Cancel	F4=List
Esc+5=Reset	Esc+6=Command	Esc+7>Edit	Esc+8=Image
Esc+9=Shell	Esc+0=Exit	Enter=Do	

13. Navigate through the licenses. When you have finished reading, click F3 twice. You are returning to the Software License Agreements panel. Select Accept License Agreements and click Enter; see Example 6-8.

Example 6-8 License agreement menu

Software License Agreements

Move cursor to desired item and press Enter.

Show Installed License Agreements
Accept License Agreements

F1=Help Esc+9=Shell	F2=Refresh Esc+0=Exit	F3=Cancel Enter=Do	Esc+8=Image
------------------------	--------------------------	-----------------------	-------------

14. Click Tab followed by Enter to accept the license and change the *no* to a *yes*, as shown in Example 6-9.

Example 6-9 Accept license agreements

Accept License Agreements

Type or select values in entry fields.
Press Enter AFTER making all desired changes.

ACCEPT Installed License Agreements [Entry Fields]
yes +

F1=Help Esc+5=Reset Esc+9=Shell	F2=Refresh Esc+6=Command Esc+0=Exit	F3=Cancel Esc+7>Edit Enter=Do	F4=List Esc+8=Image
---------------------------------------	---	-------------------------------------	------------------------

15. After the status of the command has changed to OK, click F10 or ESC+0 to leave the license agreement. The keys you need depend on the terminal you are using; see Example 6-10.

Example 6-10 Command status

COMMAND STATUS

Command: OK stdout: no stderr: no

Before command completion, additional instructions may appear below.

F1=Help Esc+8=Image	F2=Refresh Esc+9=Shell	F3=Cancel Esc+0=Exit	Esc+6=Command /=Find
------------------------	---------------------------	-------------------------	-------------------------

n=Find Next

16. The installation assistant will guide you through the first administrative tasks, such as setting a root password or configuring the network connection. Proceed with the setup as described in the AIX documentation. To complete this task and get to a login prompt, use ESC+0 or F10. You may start this installation assistant at any time again by using the command **install_assist** after login as *root*. The installation assistant is shown in Example 6-11.

Example 6-11 AIX installation assistant

Installation Assistant

Move cursor to desired item and press Enter.

Set Date and Time
Set root Password
Configure Network Communications
Install Software Applications
Using SMIT (information only)
Tasks Completed - Exit to Login

F1=Help
Esc+9=Shell

F2=Refresh
Esc+0=Exit

F3=Cancel
Enter=Do

Esc+8=Image

The installation of AIX 6.1 in the logical partition has completed. You may now start with configuring your AIX or installing your applications.



IBM i V6.1 installation

This chapter explains the installation process of the IBM i V6.1 Operating System on an IBM BladeCenter JS23/JS43 Express server installed in a BladeCenter S chassis using the disks provided in the disk storage modules.

For the IBM BladeCenter JS23/JS43 in a BladeCenter H chassis, the installation process is similar to the information provided here, except that the storage is provided from a SAN environment.

For a technical overview and complete information, as well as latest updates on IBM i on Power blades, refer to the readme file available at the following address:

<http://www-03.ibm.com/systems/power/hardware/blades/ibmi.html>

This chapter discusses the following base topics:

- ▶ “Preparing for installation” on page 242
- ▶ “IBM System Access for Windows V6R1” on page 260
- ▶ “Creating an IBM i 6.1 partition” on page 271
- ▶ “Installing the IBM i 6.1 Operating System” on page 303
- ▶ “IBM i 6.1 Backup/Restore” on page 314

7.1 Preparing for installation

There are important considerations for setting up and using IBM i 6.1 client logical partitions on IBM Power servers or the IBM BladeCenter JS23 or JS43 Express server. On Power blades, you use the Integrated Virtualization Manager (IVM) to manage partitions.

A *client logical partition* is a partition that uses some of the I/O resources of another partition.

When the IBM i 6.1 client logical partition is managed by Integrated Virtualization Manager (IVM), you can assign only virtual resources to the 6.1 partition. Disk units, optical devices, and Ethernet are accessed using virtual I/O adapters. The Virtual I/O Server (VIOS) logical partition provides the disk, optical, and network resources to the client logical partition. This configuration is sometimes referred to as a *pure virtual partition*.

A major benefit of using virtual I/O is that you can share the hardware among the client logical partitions through the server logical partition. This allows you to optimize the amount of hardware used by the host partition.

7.1.1 Software installation process

The IBM i 6.1 installation process involves three phases:

- ▶ Pre-installation activities
- ▶ Main installation
- ▶ Post-installation activities

Figure 7-1 on page 243 illustrates one type of software installation of the IBM i 6.1 Operating System. The example this process illustrates is one of performing an upgrade to a new release of the operating system. However, many of the steps are similar when performing a new installation.

Note: IBM i V6R1 can be ordered pre-installed on the JS23/JS43. Use feature code FC 8141 (IBM i pre load) and FC 8147 (VIOS Pre-load pre-req) when ordering.

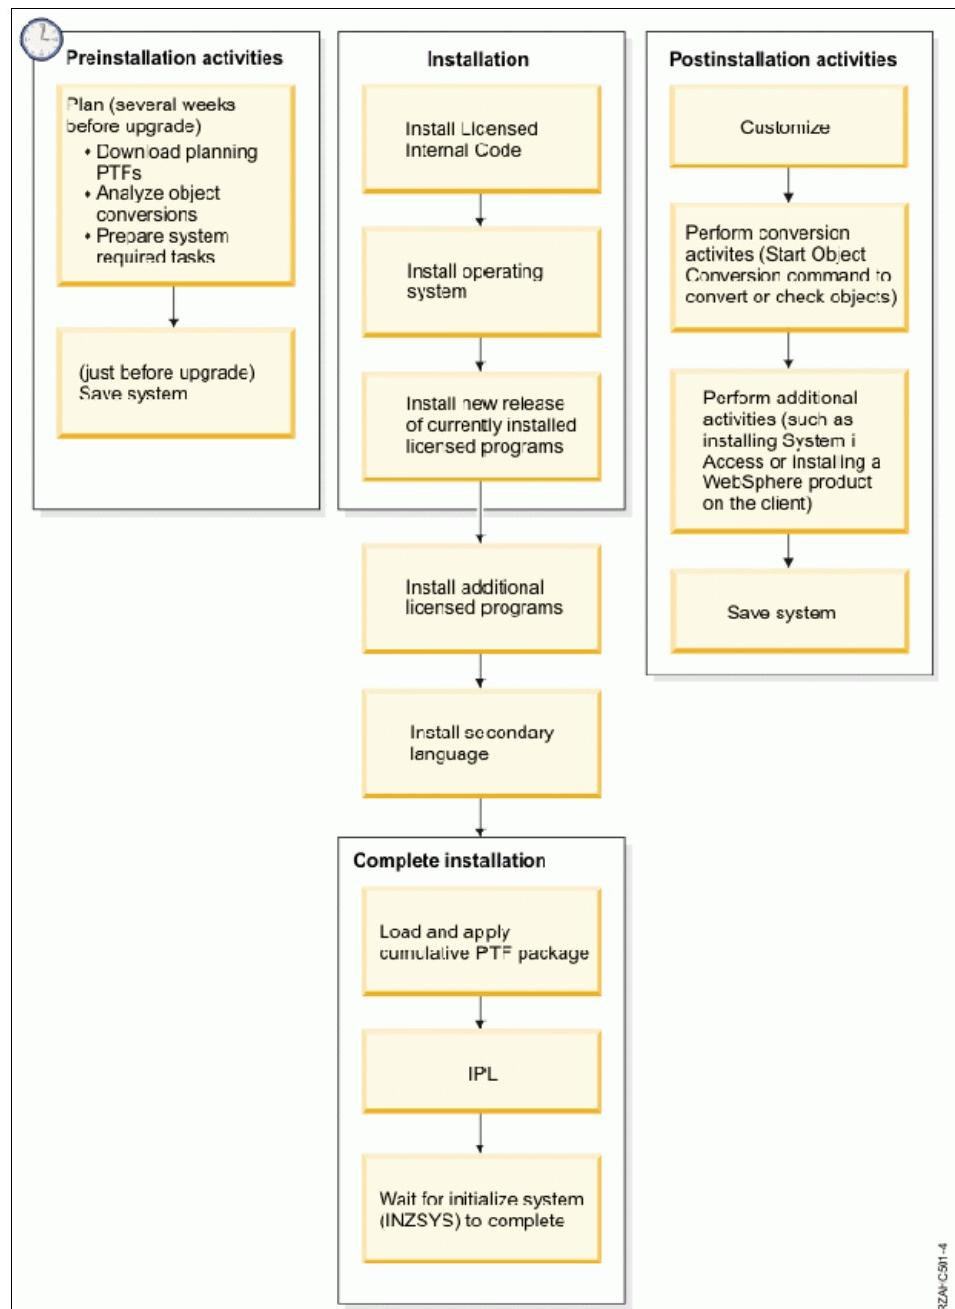


Figure 7-1 IBM i 6.1 installation process

RZAN-C501-4

7.1.2 Hardware environments

This section describes an example IBM BladeCenter chassis and IBM BladeCenter JS23/JS43 Express server configuration with recommended firmware levels.

Note: The disk configurations are dependent on the I/O requirements. For example, two SAS disk drives will not be enough with mirroring and backup to the media library. For performance reasons it is recommended to install IBM i to disk units other than the internal disks of the JS23/JS43.

A combination of BladeCenter S chassis and IBM BladeCenter JS23 or IBM BladeCenter JS43 Express server are included in the following list:

BladeCenter S chassis

- ▶ 1 Advanced Management Module
- ▶ 1 supported Ethernet I/O Module
- ▶ 1 IBM BladeCenter S 6-Disk Storage Module
- ▶ 2 SAS hot-swap disks
- ▶ 1 SAS Connectivity Module

IBM BladeCenter JS23 Express

- ▶ 1 JS23 Express server
- ▶ 4 GB memory
- ▶ SAS Expansion Card (CFFv) for IBM BladeCenter
- ▶ 1 SAS disk drive

IBM BladeCenter JS43 Express

- ▶ 1 JS43 Express server
- ▶ 4 GB memory
- ▶ SAS Expansion Card (CFFv) for IBM BladeCenter
- ▶ 1 SAS disk drive

A combination of BladeCenter H chassis and IBM BladeCenter JS23 or IBM BladeCenter JS43 Express server are included in the following list:

BladeCenter H chassis

- ▶ 1 Advanced Management Module
- ▶ 1 supported SAN I/O Module
- ▶ 1 supported Ethernet I/O Module
- ▶ 1 Multi-Switch Interconnect Module

IBM BladeCenter JS23 Express

- ▶ 1 JS23 Express server
- ▶ 4 GB memory
- ▶ 1 QLogic Ethernet and 4 GB Fibre Channel Expansion Card (CFFh)
- ▶ 1 SAS disk drive

BM BladeCenter JS43 Express

- ▶ 1 JS43 Express server
- ▶ 4 GB memory
- ▶ 1 QLogic Ethernet and 4 GB Fibre Channel Expansion Card (CFFh)
- ▶ 1 SAS disk drive

Table 7-1 lists the minimum and required features required to manage an IBM BladeCenter JS23/ Express system with the IBM i 6.1 Operating System.

Table 7-1 BladeCenter JS23 feature list

Feature	Description	Notes
7778-23X	IBM BladeCenter JS23	4-Way Single wide blade
8395	Processor Entitlement List Priced	Four processor Entitlement required
8393	Processor Entitlement Express Priced	
8229	4 GB DDE2 667 MHz VLP RDIMMs	4 GB Minimum required - up to 32 GB Maximum
8239	8 GB DDE2 667 MHz VLP RDIMMs	
8245	16 GB DDE2 400 MHz VLP RDIMMs	
8237	73 GB SAS 10k SFF disk	Disk is not required on base offering
8236	146 GB SAS 10k SFF disk	
8274	300 GB SAS 10K SFF disk	
8273	69 GB SDD SFF disk	
8252	QLogic Ethernet and 4 GB Fibre Channel Expansion Card (CFFh)	- Option for SAN connection and Ethernet in BladeCenter H - Connects to MSIM in high speed switch bays 7-8 or 9-10 - Not supported in BladeCenter S
8240	Emulex 8 GB Fibre Channel Expansion card (CIOv)	Option for SAN connection in Bay 3 or 4 of an H or S chassis

Feature	Description	Notes
8241	Qlogic 4 GB Fibre Channel Expansion card (CIOv)	Option for SAN connection in Bay 3 or 4 of an H or S chassis
8242	Qlogic 8 GB Fibre Channel Expansion card (CIOv)	Option for SAN connection in Bay 3 or 4 of an H or S chassis
8271	Qlogic 8 GB Fibre Channel Expansion card (CFFh)	

Table 7-2 on page 246 lists the minimum and required features needed to manage an IBM BladeCenter JS43 Express system with the IBM i 6.1 Operating System.

Table 7-2 BladeCenter JS43 feature list

Feature	Description	Notes
7778-23X Plus 8446	IBM BladeCenter JS43 Express Double wide 8-way blade	Adding feature 8446 to the 7778-23X equals the JS43 double wide blade package
8395	Processor Entitlement List Priced	Eight processor Entitlement required
8393	Processor Entitlement Express Priced	
8229	4 GB DDE2 667 MHz VLP RDIMMs	4 GB Minimum required - up to 32 GB Maximum
8239	8 GB DDE2 667 MHz VLP RDIMMs	
8245	16 GB DDE2 400 MHz VLP RDIMMs	
8237	73 GB SAS 10k SFF disk	Disk is not required on base offering
8236	146 GB SAS 10k SFF disk	
8274	300 GB SAS 10K SFF disk	
8273	69 GB SDD SFF disk	
8252	QLogic Ethernet and 4 GB Fibre Channel Expansion Card (CFFh)	- Option for SAN connection and Ethernet in BladeCenter H - Connects to MSIM in high speed switch bays 7-8 or 9-10 - Not supported in BladeCenter S
8240	Emulex 8 GB Fibre Channel Expansion card (CIOv)	Option for SAN connection in Bay 3 or 4 of an H or S chassis

Feature	Description	Notes
8241	Qlogic 4 GB Fibre Channel Expansion card (CIOv)	Option for SAN connection in Bay 3 or 4 of an H or S chassis
8242	Qlogic 8 GB Fibre Channel Expansion card (CIOv)	Option for SAN connection in Bay 3 or 4 of an H or S chassis
8271	Qlogic 8 GB Fibre Channel Expansion card (CFFh)	

For more information on supported devices on a BladeCenter JS23/JS43 server, refer to the following site:

<http://www.ibm.com/systems/power/hardware/blades/ibmi.html>

This site contains links to a specifications and a readme file that has excellent detail for implementing this type of configuration. There is also information relating to the different BladeCenter chassis configurations and requirements.

7.1.3 BladeCenter hardware preparation

For a smooth installation process of IBM i 6.1, all installed hardware components should be at the latest firmware levels.

The following components should be checked and updated if necessary:

- ▶ IBM JS23/JS43 Express server firmware (located on the service processor, and updated through the VIOS)
- ▶ Advanced Management Module (AMM) firmware
- ▶ Fibre Channel I/O module firmware
- ▶ Ethernet I/O module firmware
- ▶ SAS Connectivity Module firmware
- ▶ Expansion card firmware (examples: QLogic ethernet and 4 GB Fibre

For firmware updates for the IBM BladeCenter JS23 or IBM BladeCenter JS43 Express, BladeCenter S, and BladeCenter H chassis and integrated components, refer to:

<http://www.ibm.com/support/us/en/>

When you reach that window, follow these steps:

1. Click the down arrow in the Product family box and select the corresponding product: BladeCenter. Then click the blue arrow.

2. Click the down arrow in the Product family box and select the corresponding product: IBM BladeCenter JS23, BladeCenter JS43, BladeCenter S, or BladeCenter H.
3. Click the down arrow button in the Operating system box and select **IBM i 6.1**, as shown in Figure 7-2. Then click the **Go** button to activate the search.

Figure 7-2 on page 248 provides an example of the search options when using the support web site to locate updates.

The screenshot shows a search interface for IBM BladeCenter. At the top, there's a breadcrumb navigation: IBM Systems > Systems support > BladeCenter >. Below it is the title "Support for IBM BladeCenter". There are two small images of blade server racks. The main area contains four dropdown menus for filtering search results:

- Select your product**: A note below says "Fields marked with an asterisk (*) are required."
- Product family:** * (marked with an asterisk) - The dropdown menu shows "BladeCenter JS22".
- Type:** - The dropdown menu shows "All types".
- Model:** - The dropdown menu shows "All".
- Operating system:** - The dropdown menu shows "IBM i 6.1".

At the bottom right of the form is a blue "Go" button.

Figure 7-2 Firmware information and download

4. On the next window, scroll down to view all the available updates and drivers related to the selected product and operating system for downloading.

Figure 7-3 on page 250 shows an example of the available firmware and bios updates. Scroll the list to find the update you need or tailor the results using the Refine results option.

→ Alternate download resources view

Refine results: All categories

Observe all prerequisites and co-requisites listed in each download.
Click the category below to quickly jump to the code you need, then click the version to access the download page.

• Critical update	• Hard drive
• Advanced Management Module	• Infiniband
• BIOS	• Management Module
• Fibre Channel driver	• RAID
• Fibre Channel firmware	• Security
• Fibre Channel solution	• Serial attached SCSI (SAS)
• Fibre Channel utility	• Utility
• Firmware update	

→ Having trouble downloading a file from our site?

Critical update	
Frequently Asked Questions (FAQ) Seagate OEM hard drive issue - IBM BladeCenter and System x	13 Mar 2009
Advanced Management Module	
Advanced Management Module (AMM) firmware release matrix - IBM BladeCenter	30 Mar 2009
(BOFM enabled) (Functional release) Advanced Management Module firmware v2.48D (BPET48D) - IBM BladeCenter	30 Mar 2009 v2.48D
Enhanced Role Based Security Snap-in (LDAP snap-in) - IBM BladeCenter	25 Jun 2008 v2.01
BIOS	
BIOS/Firmware Inband update v3.00, 01EA340_039_039 - IBM BladeCenter JS12, JS22	21 Nov 2008 v01EA340_039_039
Firmware v01EA320_046_030 - IBM BladeCenter JS22	29 May 2008 v01EA320_046_030
Fibre Channel driver	
Boot Code for Emulex 4Gb fibre channel expansion card v6.02a7 - IBM BladeCenter	29 May 2008 v6.02a7

Figure 7-3 Example: Partial list of available downloads by type

7.1.4 VIO Server software environments

VIO Server is part of IBM PowerVM Editions (formerly Advanced POWER Virtualization). It is required in the IBM i 6.1 for IBM BladeCenter JS23/JS43 Express environment. At minimum, VIOS level 1.5 is required for IBM i. It is recommended to use version 2.1 or later.

Work with your local sales channel to ensure that PowerVM (Standard or Enterprise Edition) and the latest fix pack are part of the BladeCenter JS23/JS43 order. Consult the **Supported Environments** page to verify that you have the minimum supported release of VIOS:

<http://www.ibm.com/systems/power/hardware/blades/ibmi.html>

For detailed information about using the Integrated Virtualization Manager and VIO Server, refer to Chapter 4, “System planning and configuration using VIOS with IVM” on page 71.

7.1.5 Network considerations

Before the initial installation process starts, a set of IP addresses should be reserved for setting up the minimum environment.

AMM

The Advanced Management Module address is a physical LAN IP address. It is used to manage the BladeCenter chassis and IBM BladeCenter JS23/JS43 Express server remotely.

Note: We suggest that you do not place the Advanced Management Module (AMM) IP address in the same subnet as the other addresses to prevent I/O traffic between AMM and integrated Ethernet Switch modules (Proxy-Loop).

Ethernet I/O module

This IP address is used to connect the Ethernet I/O Module to the physical LAN, thus allowing any blades in the BladeCenter chassis to access the LAN.

VIOS/IVM

This IP address is used to connect the Virtual I/O Server (VIOS) and Integrated Virtualization Manager (IVM) to the LAN.

SAS Modules

This IP address is used to communicate with the SAS modules.

6.1 LAN console	This IP address on the LAN is used to allow the 5250 console to connect to the VIOS using the IBM System i Access for Windows software.
6.1 production interface	This IP address on the external LAN is used to provide 5250 production network access. This address will be configured after 6.1 is installed using LAN console. It is recommended that the 6.1 LAN console and production network interface use two separate Virtual Ethernet adapters in the 6.1 partition.
PC for LAN Console	When the System i for Windows LAN console is first established, the PC console must be on the same subnet as the 6.1 partition. After the console is established, this restriction is removed.

7.1.6 Storage consideration BladeCenter H

An IBM BladeCenter JS23/JS43 Express server installed in a BladeCenter H chassis has no access to any physical devices in the BladeCenter H chassis. Storage must be provided by attaching LUNs on a SAN to the VIO Server. The VIO Server provides a virtualized access from IBM i on the JS23/JS43 to the LUNs.

Note: When you configure LUNs for IBM i, configure them as 512-byte AIX LUNs, not as 520-bytes IBM i LUNs.

For detailed information about how to define a SAN environment for an IBM BladeCenter JS23 or JS43 blade installed in a BladeCenter H chassis using LUN attached disks to install IBM i 6.1, refer to:

http://www.ibm.com/systems/power/hardware/blades/i_on_blade_readme.pdf

7.1.7 Disk consideration in BladeCenter S

This section describes the configuration environment used in this chapter to install IBM i 6.1 on an IBM BladeCenter JS23/JS43 Express server in a BladeCenter S chassis.

IBM i 6.1 running on an IBM BladeCenter JS23/JS43 Express server does not have physical access to storage or other devices; instead, the hardware does. Every hardware device is provided to the IBM BladeCenter JS23/JS43 server as a virtual device using the Integrated Virtualization Manager (IVM) functionality.

To provide access to a SAS drive in the BladeCenter S chassis to the partition, at least one SAS I/O module must be installed in the BladeCenter S chassis. An SAS expansion adapter (CIOv) also must be installed in each IBM BladeCenter JS23 or IBM BladeCenter JS43 Express server. A single SAS I/O module provides access to both Disk Storage Modules (DSM) and all 12 disks.

The physical connection to tape drives is owned and managed by VIOS. The IBM i does not have direct access to the tape. The SAS I/O expansion module also provides a connection of an LTO tape connection for backup and restore.

After SAS disk drives are assigned to an IBM BladeCenter JS23/JS43, they become available in an hdiskX hard disk drive. Each virtualized SAS disk drive will be recognized as a DDxx physical disk drive in IBM i 6.1 after it has been installed.

Note: SATA drives installed in the Disk Storage Modules (DSM) in the BladeCenter S chassis are not supported for IBM i 6.1.

The concept of a chassis with a self-contained disk subsystem or storage modules is unique to the BladeCenter S chassis. The module is a collection of disk drives that are made accessible through a SAS module and a SAS card in the BladeCenter JS23/JS43 server. The SAS module is responsible for both the provisioning of physical disk drives through zoning and for failover redundancy when two SAS modules are present. Each SAS drive in the DSM is assigned individually.

Zoning allows you to map hard drives in storage module 1 and storage module 2 to the blade servers, and to map the blade servers to the external ports on the connectivity module.

Note: To provide better RAS, use the IBM i 6.1 disk mirroring functionality for the SAS disks provided from the VIO Server; this will protect the IBM i 6.1 installation.

To ensure higher availability, the configuration of assigned SAS disk drives in the Disk Storage Modules (DSM) can be individually configured. Information provided in 7.1.8, “Disk configuration in BladeCenter S” on page 254, helps you to visualize or select a predefined DSM configuration.

7.1.8 Disk configuration in BladeCenter S

To use a pre-defined configuration to a BladeCenter JS23/JS43 server, you must establish a connection to the SAS Module as shown in Figure 7-4 on page 254. using a browser window directly connected to the SAS Module.

An alternative, that is more intuitive for clients is the SCM GUI.

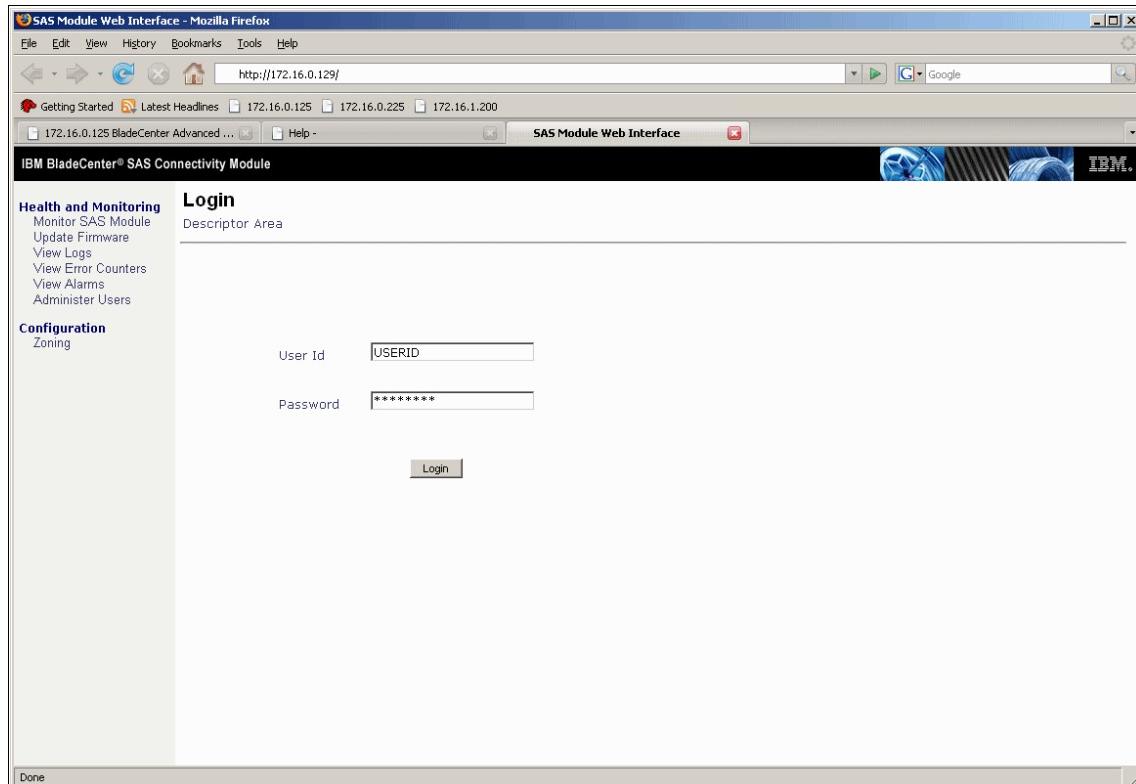


Figure 7-4 SAS Connection module login

1. Enter **User ID** and **Password** of the account that has access to the SAS module and click **Login**. An example of the next screen that appears in the SAS module welcome screen is provided in Figure 7-5 on page 255.

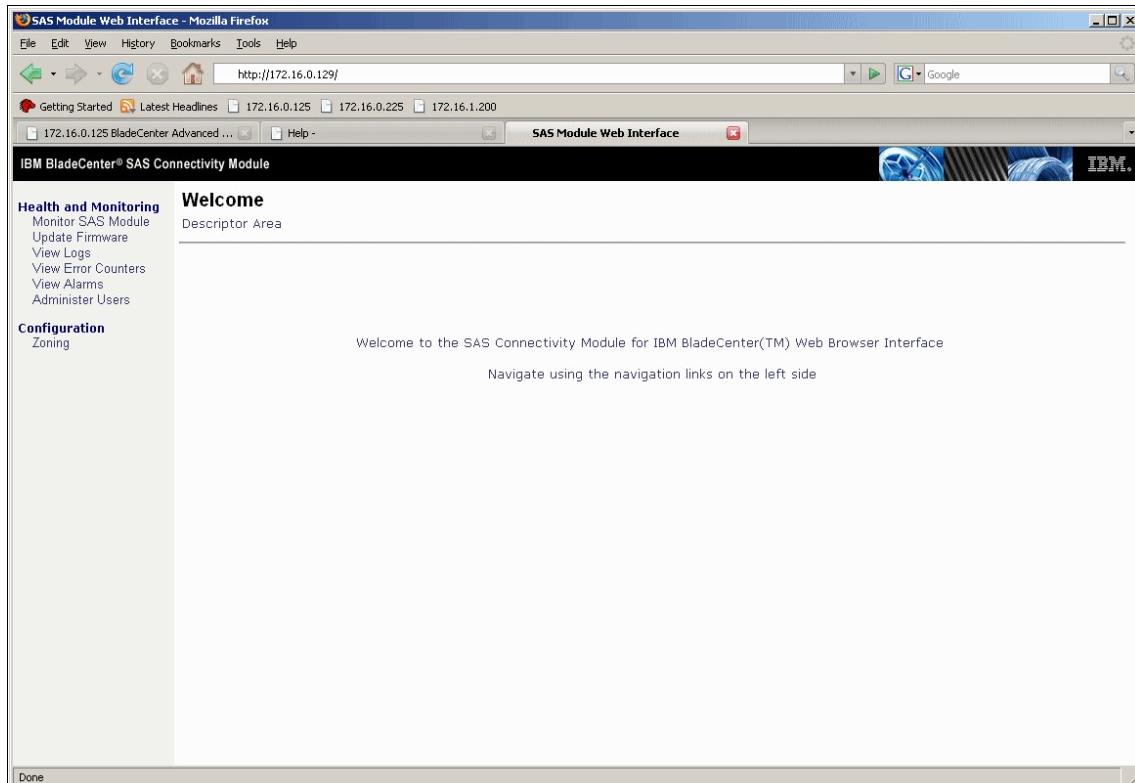


Figure 7-5 SAS connection module welcome

2. Select **Zoning**. In the example shown in Figure 7-6 on page 256, **Predefined Config 09** is selected and active. Notice that our BladeCenter JS23/JS43 installed in slot 4, Zone Group ID 37 is configured. Remember the Zone Group ID for the following window to examine the corresponding hard disk drives.
3. Click **Basic Zone Permission Table**.

The screenshot shows the 'Manage Fabric' section of the SAS Module Web Interface. On the left, there's a sidebar with 'Health and Monitoring' and 'Configuration' (Zoning selected). The main area has a 'Descriptor Area' header. Below it, two configurations are listed: 'Working Configuration' (selected) and 'Active Configuration'. A dropdown menu shows 'Predefined Config 09' (selected) and 'Predefined Config 09'. A button 'Activate this Configuration' is present. The central part is a table titled 'Basic Zone Permission Table' under 'Zone Groups'. The table has columns for Zone Group ID, External, Blade, Storage Module, SAS Module 1 Ports, Disk 1, SAS Module 2 Ports, and Disk 2. Rows show various zone group IDs (e.g., 1, 30, 31, 32, 33, 34, 35, 36, 37, 38, 39, 61, 62, 63, 64, 65, 66) and their corresponding port assignments. Row 37 is circled in red.

Zone Group ID	External	Blade	Storage Module	SAS Module 1 Ports	Disk 1	SAS Module 2 Ports	Disk 2
1				1,2	1		1
30	1						
31	2						
32	3						
33	4						
34			1				
35			2				
36			3				
37			4				
38			5				
39			6				
61					1		
62					2		
63					3		
64					4		
65					5		
66					6		

Figure 7-6 SAS connection module zone groups

Figure 7-7 on page 257 shows the definition and setup window for the actual configuration. In this configuration three disks from SAS module 1 and three disks from SAS module 2 are defined for Predefined Config 09. Individual User Defined Configs are provided for specific configurations.

For more detailed information about this topic, refer to: *Implementing the IBM BladeCenter S Chassis*, SG24-76827.

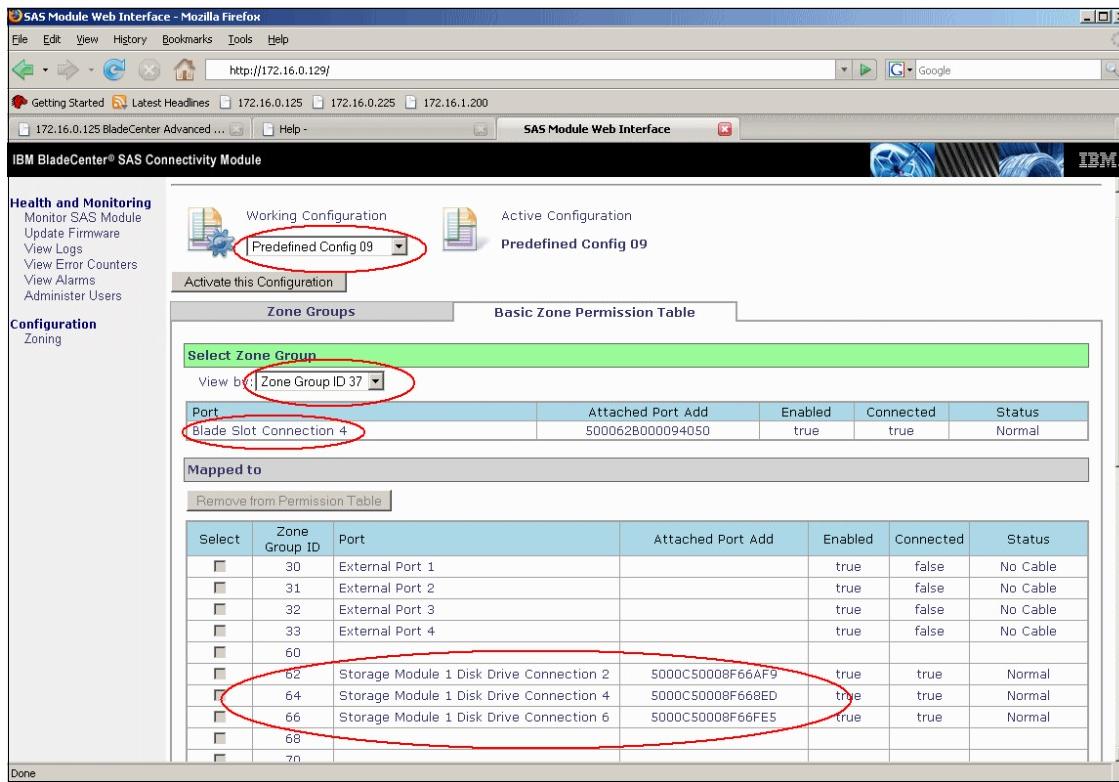


Figure 7-7 SAS connection module zoning

4. To verify the configuration in the SAS module configuration menus, logon to the IBM BladeCenter Advanced Management Module. Under Storage Tasks, select **Configuration** as shown in Figure 7-8 on page 258.

Zone Configuration Management for I/O Modules

The table below displays zone configurations stored on the given I/O Module. Please select the desired zone configuration from the list and activate it. The Refresh button would be helpful in refreshing the status once the zone is applied. If you have multiple SAS Connectivity Modules installed and both are in working order, a check box will be provided that allows you to easily apply the same configuration to each I/O Module. The default setting is to apply the same zone configuration to each. If you uncheck the check box, information for both I/O Modules will be presented and you can select a zone configuration from each. However, it is highly recommended that you select the same zone configuration for both I/O Modules.

Select	Active?	Name	Type	Intended # of Blades	Intended # of SAS Modules	Max Disks per Blade	Configuration Store	Date
<input checked="" type="radio"/>		Predefined Config 03	Pre-defined	6	1	2	7	04/24/2007, 02:00:00
<input type="radio"/>	<input checked="" type="checkbox"/>	Predefined Config 09	Pre-defined	2	1	6	13	04/24/2007, 02:00:00

Figure 7-8 AMM SAS configuration zone

5. Click **Predefined Config 09** to proceed.

Figure 7-9 on page 259 shows the current configuration. Select the blade in the upper rectangle to highlight the assigned disks to that blade.

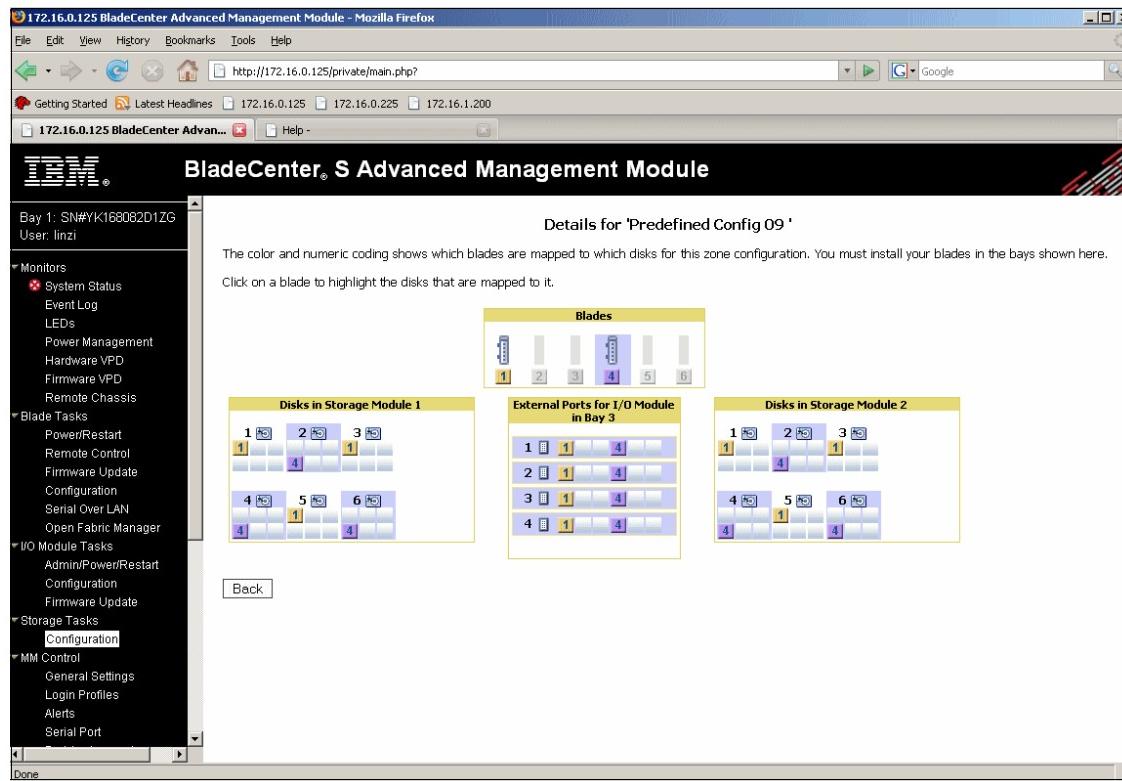


Figure 7-9 AMM SAS configuration zone 9

For detailed information, refer to *Implementing the IBM BladeCenter S Chassis*, SG24-76827 and *IBM BladeCenter Products and Technology*, SG24-7523.

7.1.9 Individual BladeCenter S disk configuration

If one of the eight predefined SAS I/O module disk configurations does not match the target configuration, four user-predefined configurations are available for individual use.

To create an individual customer configuration, one of the four user predefined configuration templates must be changed. Two access methods are available to create a custom configuration:

- ▶ Using Telnet functionality to access the SAS I/O module command line interface
- ▶ Using the Storage Configuration Manager (SCM) graphical interface

The IBM Storage Configuration Manager (SCM) may be used to create an individual configuration if you are not familiar using the SAS I/O module command line interface. The SCM software can be downloaded from:

<https://www-304.ibm.com/systems/support/supportsite.wss/docdisplay?lnid=cid=MIGR-5502070&brandind=5000008>

7.2 IBM System Access for Windows V6R1

System i Access for Windows fully integrates the power of the IBM i 6.1 operating system into the desktop so that users have a simple and optimized interface to use the systems capabilities in business environments. The product provides a single solution to work with databases and other data on the system, to run 5250 applications, and to administer the IBM i environment.

The main reason for installing this software is to access the console session on your IBM i partition. System i Access provides two different connection options for the system console. One option is a direct cable from an adapter to a PC comm port and the other method is to connect using the LAN. When you install IBM i in the blade server environment you can only use the LAN Console option for the partition. There are no physical adapters compatible with the blade environment that can be used for a direct console connection.

Prior to 6.1, the IBM System i Access for Windows product was known as the IBM iSeries® Access for Windows products.

With System i Access for Windows, a full suite of client/server capabilities to support a PC and system interaction:

- ▶ System i Access for Windows required programs
- ▶ System i Access for Windows optional features
- ▶ System i Navigator graphical interface and functions
- ▶ Data access functions
- ▶ AFP Workbench Viewer
- ▶ Secure Socket Layer (SSL)
- ▶ Printer Driver
- ▶ Display and Printer Emulation
- ▶ Operations Console

7.2.1 Preparation System i Access for Windows

To meet all the System i Access for Windows hardware requirements, follow the instructions described in:

<http://www-03.ibm.com/systems/i/software/access/windows/v6r1pcreq.html>

For more information about the IBM System i Access for Windows V6R1, see:

<http://www.ibm.com/systems/i/software/access/index.html>

To obtain the IBM System i Access for Windows software, go to the following address:

<http://www.ibm.com/systems/i/software/access/caorder.html>

Note: When the IBM i Access for Windows connection is first established, the console PC must be on the same subnet as the 6.1 partition. After the console is established, this restriction is removed.

7.2.2 Installation System i Access for Windows

This section describes the sequence to install System i Access for Windows.

1. After you insert the System i Access for Windows V6R1 DVD, the window shown in Figure 7-10 appears. Select the language you require, and click **OK** to continue.

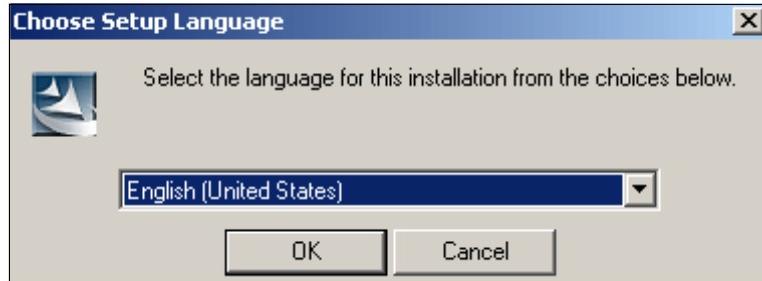


Figure 7-10 Setup LAN Console

2. The InstallShield Wizard Welcome window shown in Figure 7-11 on page 262 appears. Click **Next** to continue.



Figure 7-11 IBM System i Access for Windows welcome screen

3. The License Agreement shown in Figure 7-12 on page 263 appears. You can select **I accept the terms in the license agreement**. Click **Next** to continue.

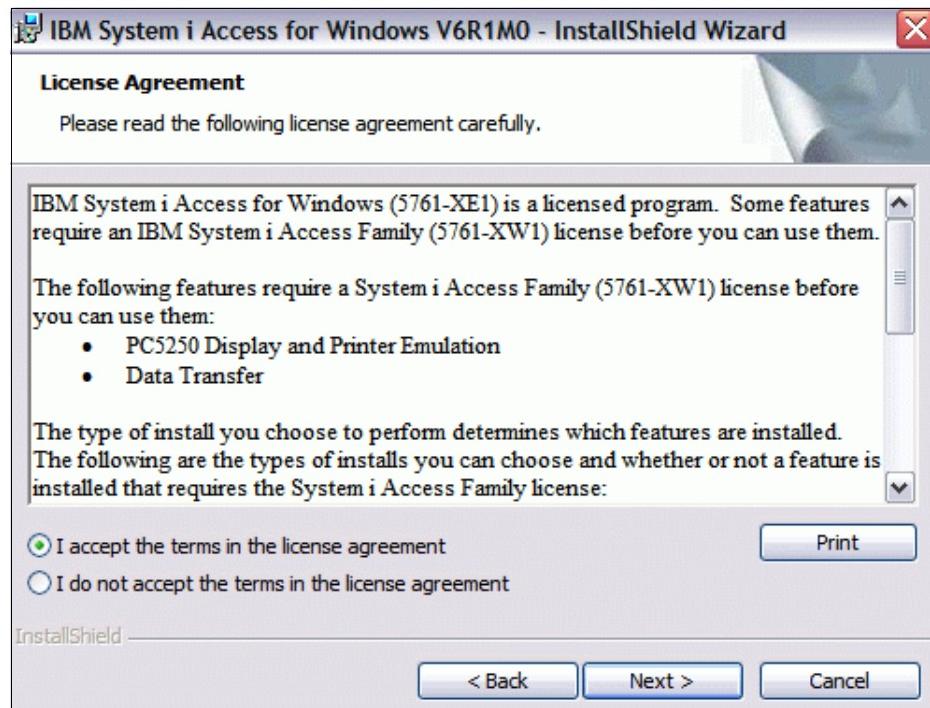


Figure 7-12 IBM System i Access for Windows Licence Agreement

4. IBM System i Access for Windows can be installed at a different location, as shown in Figure 7-13 on page 264.

To store the software at a different location, click **Change...** and choose a new location.

Or, accept the predefined path and click **Next** to continue.



Figure 7-13 IBM System i Access for Windows install location

5. Depending on the native language, a selection can be made in the following window as shown in Figure 7-14 on page 265. Normally the same language will be chosen as the language for the IBM i 6.1 Operating System.

Click **Next** to continue.



Figure 7-14 IBM System i Access for Windows Primary language

6. Depending on the complexity of functions, several choices are available as shown in Figure 7-15 on page 266. The normal case is a complete installation. Experienced administrators can select the custom installation to save disk space, or install determined functions only.

Click **Next** to continue.

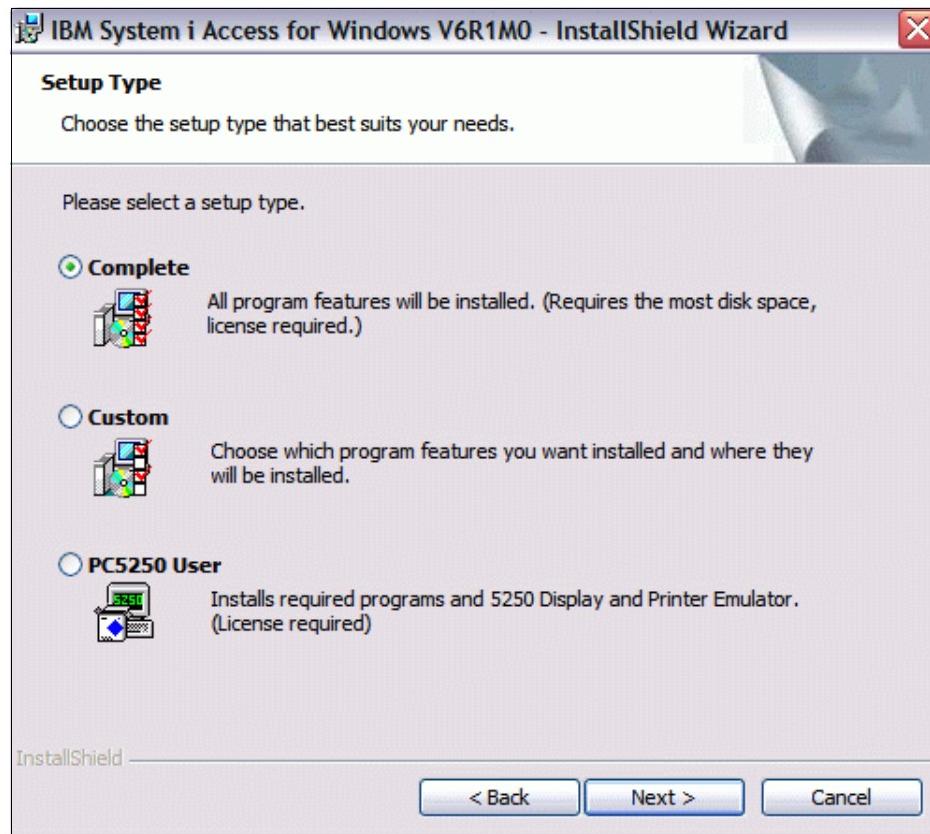


Figure 7-15 IBM System i Access for Windows Setup Type

7. Select **Complete** and click **Next**
8. Some features require a license agreement to use their functionality, as shown in Figure 7-16 on page 267. Ask your service representative to receive a valid license key.
Click **Next** to continue.

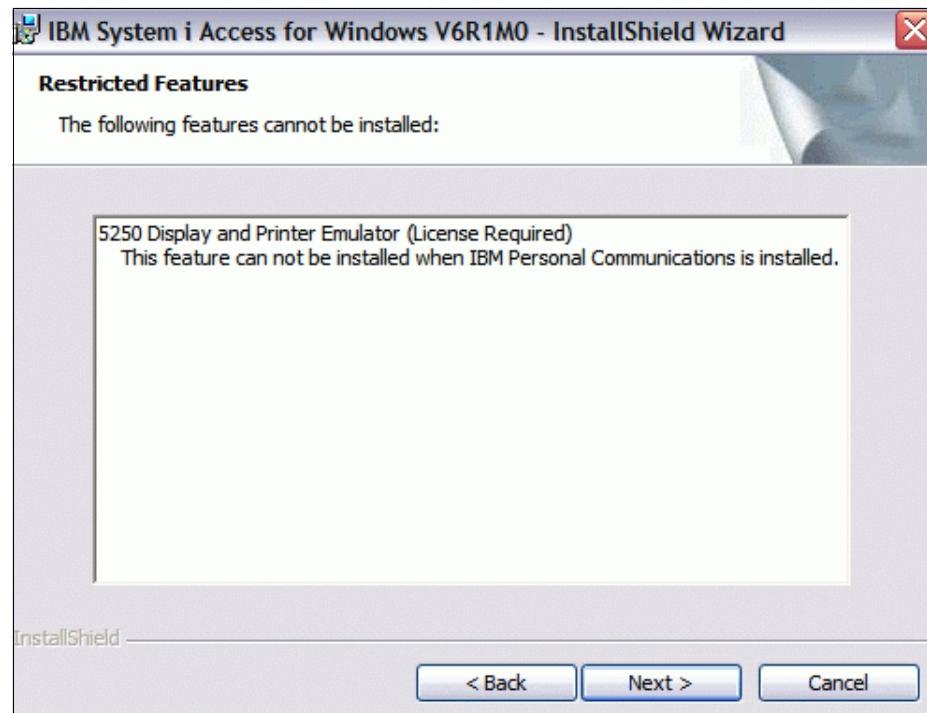


Figure 7-16 IBM System i Access for Windows Restricted Features

9. The installation starts automatically after you select **Next** in the previous menu. Figure 7-17 on page 268 shows the progress of the installation process.



Figure 7-17 IBM System i Access for Windows installation progress

10. Figure 7-18 on page 269 indicates the installation process was successful.
Click **Finish** to continue.



Figure 7-18 IBM System i Access for Windows installation completed

11. To finalize the IBM i Access for Windows installation a reboot is required, as indicated in Figure 7-19.

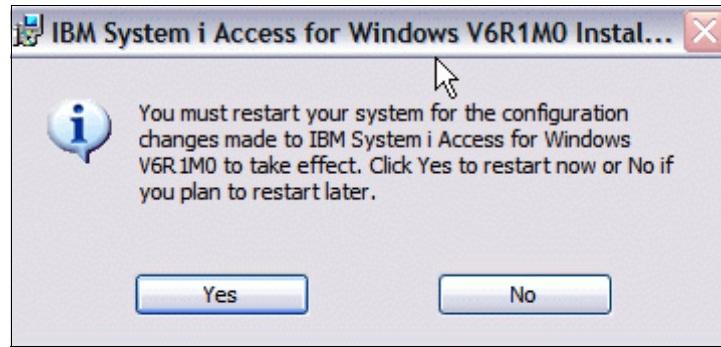


Figure 7-19 IBM System i Access for Windows Reboot

12. Click **Yes** to reboot the system.

After the console PC is successfully rebooted, the information screen shown in Figure 7-20 is displayed. The Welcome window provides additional information about the software just installed. (For some information, the administration PC needs a connection to the Internet.)



Figure 7-20 IBM System i Access for Windows Welcome

After the System i Access for Windows product is installed, you must meet two conditions before you can run a System i Access for Windows application:

- ▶ Your PC and the system must be on a TCP/IP network.
- ▶ A user ID must exist on the IBM i 6.1 Operating System.

7.3 Creating an IBM i 6.1 partition

Using Integrated Virtualization Manager (IVM) to create an IBM i 6.1 partition is similar to using the HMC. IVM uses a number of defaults that simplify partition creation. For example, because IBM i 6.1 partitions cannot own physical hardware on an IVM-managed system such as a BladeCenter JS23/JS43, those screens are omitted from the creation wizard. Other screens are simplified as well that relate to shared processor pool settings and memory settings. Typically you will need to tailor the partition properties to change the settings so they are more granular.

One other difference between the HMC method and using IVM, if you are familiar at all with the HMC you know that the partition has properties and the partition profile contains the definition of your selections for hardware, load source, console etc. With the HMC you may even have multiple partition profiles for one partition. In IVM, there is only one entity that describes both the partition properties and the partition profile. There are not separate partition profiles in IVM.

Since your IBM i partition is virtual, IVM defaults the load source and alternate IPL adapters to the Virtual SCSI client adapter in the IBM i 6.1 partition, and the console adapter to the first Virtual Ethernet adapter. If you plan to use separate Virtual Ethernet adapters for LAN console and production traffic, and you want to use the second Virtual Ethernet adapter for the LAN console, you can make the change in the partition properties.

7.3.1 IBM i 6.1 minimum requirements

The minimum recommended amount of memory for an IBM i 6.1 client partition on the BladeCenter JS23/JS43 blade is 1 GB. If you running two or more IBM i 6.1 partitions on the same blade, ensure there is enough memory on your blade.

The actual memory and CPU values should be sized individually for each IBM i 6.1 workload using the IBM Workload Estimator, which is available at:

<http://www.ibm.com/systems/support/tools/estimator/index.html>

Note: Also consider the system firmware or POWER Hypervisor for the BladeCenter JS23/JS43 blade will allocate some memory for its usage. This amount will vary but is generally less than 1GB. This value can be found on the main panel of IVM listed by *Reserved firmware memory*.

7.3.2 VIO Server configuration

For a detailed explanation of how to setup and configure the VIOS partition to use the Integrated Virtualization Manager (IVM), refer to Chapter 4, “System planning and configuration using VIOS with IVM” on page 71.

7.3.3 Creating an IBM i partition

This section provides a brief explanation of how to create an IBM i 6.1 partition.

It is assumed you have previously configured disk space (LUNs) for this partitions usage.

To create an IBM i 6.1 partition using the Integrated Virtualization Manager (IVM), follow these steps:

1. Click **View/Modify Partitions**.

Figure 7-21 on page 272 shows an example of the View/Modify partitions screen and the Create Partition button.

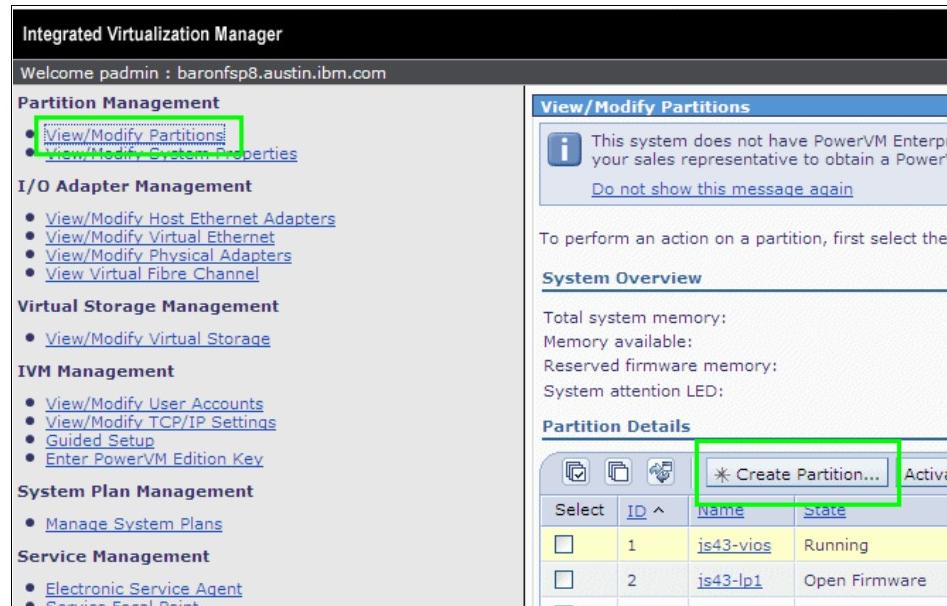


Figure 7-21 View Modify Partitions - Create Partition option

2. Click **Create Partition**. The next available Partition ID is preselected.
 - a. Enter a name for the partition in the Partition name field.
 - b. Select **IBM i** or **i5/OS** in the Environment field.
 - c. Click **Next** to continue.

Figure 7-22 on page 273 shows an example of the partition ID, name and environment fields. The ID will be filled in by the wizard using the next available number. You can change this if you desire to an unused number.

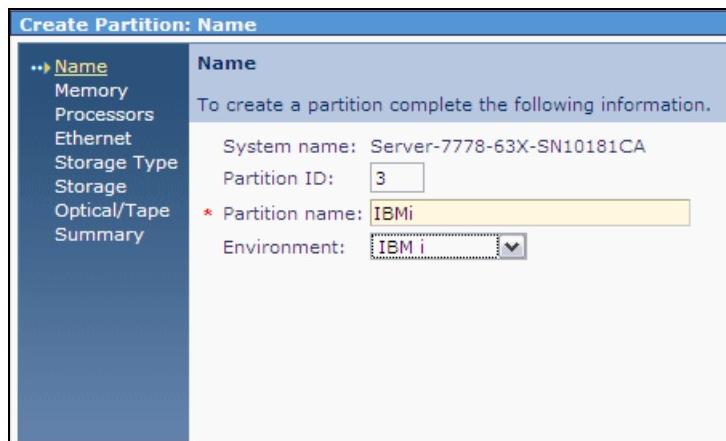


Figure 7-22 Partition id, name and environment options

3. Define the Assigned memory value (Assigned equates to Desired on the HMC). Available memory for the partition will be shown above the Assigned memory field. Click **Next** to continue.

Figure 7-23 on page 274 shows an example of the memory definitions. You can see the total system memory and what memory is available for partition usage. The available memory will be the value of the total minus allocated partition memory and hypervisor allocated memory.

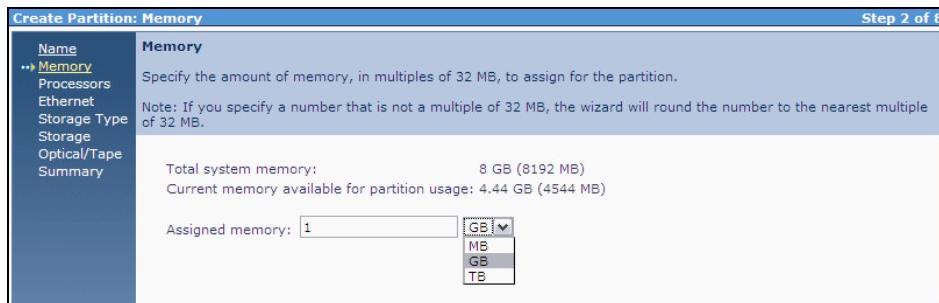


Figure 7-23 Partition memory definition panel

4. Select the desired processor configuration. Click **Next** to continue.

Figure 7-24 on page 274 is an example of the processor selection panel. In this example the blade server had 8 processors total. There are other partitions created which also use some processor capacity. In the *Assigned processors* field you will choose how many processor units to assign to this partition. For example, using shared if you choose 1 as shown in the graphic, you will have .10 units configured in the partition. The value that shows 63 available virtual processors indicates that out of the 80 (8.0 physical processors) virtual processors available, the blade has 17 virtual processors in use by other partitions. That also equates to 1.7 processors used.

Dedicated processing is another option where one or more whole processors can be assigned to the partition.

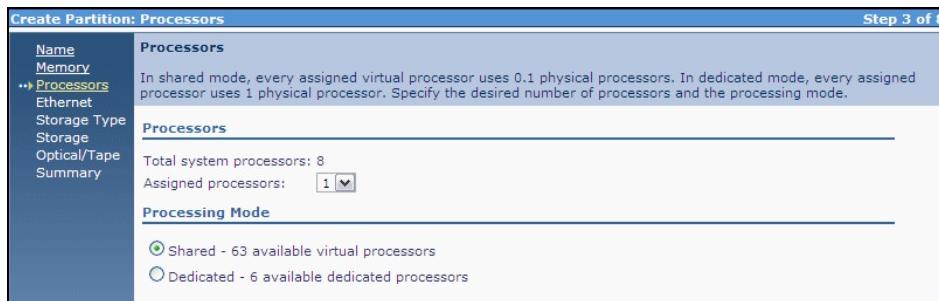


Figure 7-24 Partition processor selection

5. Specify the VLAN used for bridging.

- a. Click **Next** to continue.

Figure 7-25 on page 275 provides an example of the ethernet options available to select for this partition. Since our console for the IBM i partition is LAN based, we must have a ethernet selected that is bridged. In our example we bridged one

of the HEA ports prior to creating this partition. For more information on bridging the HEA ports see 4.5.1, “Host Ethernet Adapters” on page 100.

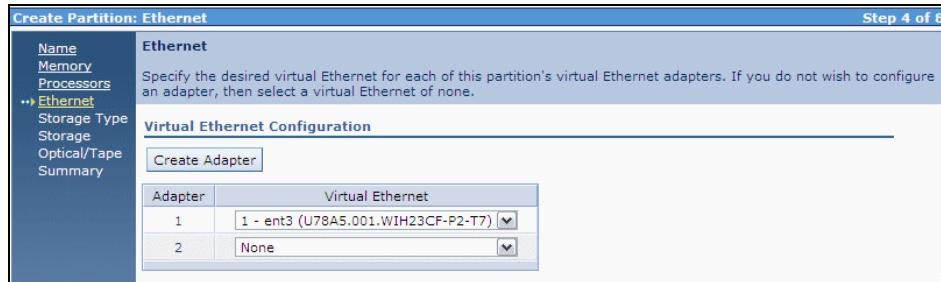


Figure 7-25 Partition ethernet selection

6. Select **Assign existing virtual disks and physical volumes**.
7. Click **Next** to continue.

Figure 7-26 on page 275 shows an example of the selection for disk units to use in the partition. You can use virtual disks or physical disks. For an IBM i partition it is recommended to use physical volumes. Selecting virtual disks will work but, you may have performance issues in the partition. It is also assumed at this point that you have previously configured LUNs or have physical units available to assign to this partition.

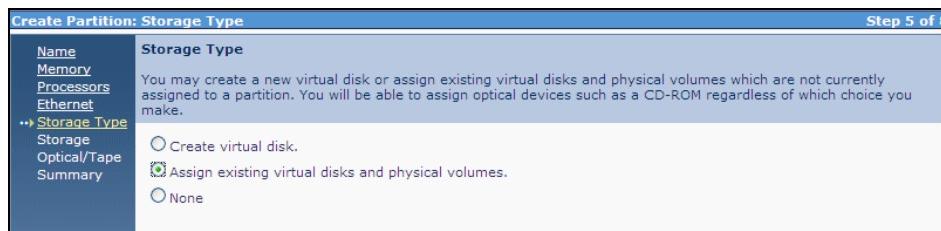


Figure 7-26 Select storage type

8. Depending on the BladeCenter chassis configuration in the storage window, you can select either:
 - a. A LUN that is already configured for the IBM i 6.1 partition from the list of **Available Physical Volumes**.
 - b. Or a disk provided from the VIO Server to the IBM i 6.1 partition from the list **Available Virtual Disks**. This disk may have been created from a storage pool.
9. Click **Next** to continue.

Figure 7-27 on page 276 shows an example of available disk units. For this partition we selected to use hdisk8 and hdisk9 which are LUNs created in a storage subsystem that have been assigned to this JS43. Depending on your configuration you may also have virtual disks available. If so, they would be listed under the *Available Virtual Disks* section.

Storage																																							
Storage																																							
Select any number of physical volumes and virtual disks from the following lists of devices which are assigned to a partition. You may use the Storage Management functions to change assignments at any time.																																							
Available Virtual Disks																																							
<table border="1"> <thead> <tr> <th>Select</th><th>Name ^</th><th>Storage Pool</th><th></th></tr> </thead> <tbody> <tr> <td><input type="checkbox"/></td><td>hdisk1</td><td>68.37 GB</td><td>U78A5.001.WIH23CF-P1-T5-L60000-L0</td></tr> <tr> <td><input type="checkbox"/></td><td>hdisk2</td><td>30 GB</td><td>U78A5.001.WIH23CF-P1-C11-L1-T1-W500507630E87FE3F-L4010400</td></tr> <tr> <td><input type="checkbox"/></td><td>hdisk3</td><td>30 GB</td><td>U78A5.001.WIH23CF-P1-C11-L1-T1-W500507630E87FE3F-L4010400</td></tr> <tr> <td><input type="checkbox"/></td><td>hdisk4</td><td>30 GB</td><td>U78A5.001.WIH23CF-P1-C11-L1-T1-W500507630E87FE3F-L4010400</td></tr> <tr> <td><input type="checkbox"/></td><td>hdisk6</td><td>20 GB</td><td>U78A5.001.WIH23CF-P1-C11-L1-T1-W500507630E87FE3F-L4010401</td></tr> <tr> <td><input type="checkbox"/></td><td>hdisk7</td><td>20 GB</td><td>U78A5.001.WIH23CF-P1-C11-L1-T1-W500507630E87FE3F-L4010401</td></tr> <tr> <td><input checked="" type="checkbox"/></td><td>hdisk8</td><td>20 GB</td><td>U78A5.001.WIH23CF-P1-C11-L1-T1-W500507630E87FE3F-L4010401</td></tr> <tr> <td><input checked="" type="checkbox"/></td><td>hdisk9</td><td>20 GB</td><td>U78A5.001.WIH23CF-P1-C11-L1-T1-W500507630E87FE3F-L4010401</td></tr> </tbody> </table>				Select	Name ^	Storage Pool		<input type="checkbox"/>	hdisk1	68.37 GB	U78A5.001.WIH23CF-P1-T5-L60000-L0	<input type="checkbox"/>	hdisk2	30 GB	U78A5.001.WIH23CF-P1-C11-L1-T1-W500507630E87FE3F-L4010400	<input type="checkbox"/>	hdisk3	30 GB	U78A5.001.WIH23CF-P1-C11-L1-T1-W500507630E87FE3F-L4010400	<input type="checkbox"/>	hdisk4	30 GB	U78A5.001.WIH23CF-P1-C11-L1-T1-W500507630E87FE3F-L4010400	<input type="checkbox"/>	hdisk6	20 GB	U78A5.001.WIH23CF-P1-C11-L1-T1-W500507630E87FE3F-L4010401	<input type="checkbox"/>	hdisk7	20 GB	U78A5.001.WIH23CF-P1-C11-L1-T1-W500507630E87FE3F-L4010401	<input checked="" type="checkbox"/>	hdisk8	20 GB	U78A5.001.WIH23CF-P1-C11-L1-T1-W500507630E87FE3F-L4010401	<input checked="" type="checkbox"/>	hdisk9	20 GB	U78A5.001.WIH23CF-P1-C11-L1-T1-W500507630E87FE3F-L4010401
Select	Name ^	Storage Pool																																					
<input type="checkbox"/>	hdisk1	68.37 GB	U78A5.001.WIH23CF-P1-T5-L60000-L0																																				
<input type="checkbox"/>	hdisk2	30 GB	U78A5.001.WIH23CF-P1-C11-L1-T1-W500507630E87FE3F-L4010400																																				
<input type="checkbox"/>	hdisk3	30 GB	U78A5.001.WIH23CF-P1-C11-L1-T1-W500507630E87FE3F-L4010400																																				
<input type="checkbox"/>	hdisk4	30 GB	U78A5.001.WIH23CF-P1-C11-L1-T1-W500507630E87FE3F-L4010400																																				
<input type="checkbox"/>	hdisk6	20 GB	U78A5.001.WIH23CF-P1-C11-L1-T1-W500507630E87FE3F-L4010401																																				
<input type="checkbox"/>	hdisk7	20 GB	U78A5.001.WIH23CF-P1-C11-L1-T1-W500507630E87FE3F-L4010401																																				
<input checked="" type="checkbox"/>	hdisk8	20 GB	U78A5.001.WIH23CF-P1-C11-L1-T1-W500507630E87FE3F-L4010401																																				
<input checked="" type="checkbox"/>	hdisk9	20 GB	U78A5.001.WIH23CF-P1-C11-L1-T1-W500507630E87FE3F-L4010401																																				
Available Physical Volumes																																							
<table border="1"> <thead> <tr> <th>Select</th><th>Name ^</th><th>Size</th><th>Physical Location Code</th></tr> </thead> <tbody> <tr> <td><input type="checkbox"/></td><td>hdisk1</td><td>68.37 GB</td><td>U78A5.001.WIH23CF-P1-T5-L60000-L0</td></tr> <tr> <td><input type="checkbox"/></td><td>hdisk2</td><td>30 GB</td><td>U78A5.001.WIH23CF-P1-C11-L1-T1-W500507630E87FE3F-L4010400</td></tr> <tr> <td><input type="checkbox"/></td><td>hdisk3</td><td>30 GB</td><td>U78A5.001.WIH23CF-P1-C11-L1-T1-W500507630E87FE3F-L4010400</td></tr> <tr> <td><input type="checkbox"/></td><td>hdisk4</td><td>30 GB</td><td>U78A5.001.WIH23CF-P1-C11-L1-T1-W500507630E87FE3F-L4010400</td></tr> <tr> <td><input type="checkbox"/></td><td>hdisk6</td><td>20 GB</td><td>U78A5.001.WIH23CF-P1-C11-L1-T1-W500507630E87FE3F-L4010401</td></tr> <tr> <td><input type="checkbox"/></td><td>hdisk7</td><td>20 GB</td><td>U78A5.001.WIH23CF-P1-C11-L1-T1-W500507630E87FE3F-L4010401</td></tr> <tr> <td><input checked="" type="checkbox"/></td><td>hdisk8</td><td>20 GB</td><td>U78A5.001.WIH23CF-P1-C11-L1-T1-W500507630E87FE3F-L4010401</td></tr> <tr> <td><input checked="" type="checkbox"/></td><td>hdisk9</td><td>20 GB</td><td>U78A5.001.WIH23CF-P1-C11-L1-T1-W500507630E87FE3F-L4010401</td></tr> </tbody> </table>				Select	Name ^	Size	Physical Location Code	<input type="checkbox"/>	hdisk1	68.37 GB	U78A5.001.WIH23CF-P1-T5-L60000-L0	<input type="checkbox"/>	hdisk2	30 GB	U78A5.001.WIH23CF-P1-C11-L1-T1-W500507630E87FE3F-L4010400	<input type="checkbox"/>	hdisk3	30 GB	U78A5.001.WIH23CF-P1-C11-L1-T1-W500507630E87FE3F-L4010400	<input type="checkbox"/>	hdisk4	30 GB	U78A5.001.WIH23CF-P1-C11-L1-T1-W500507630E87FE3F-L4010400	<input type="checkbox"/>	hdisk6	20 GB	U78A5.001.WIH23CF-P1-C11-L1-T1-W500507630E87FE3F-L4010401	<input type="checkbox"/>	hdisk7	20 GB	U78A5.001.WIH23CF-P1-C11-L1-T1-W500507630E87FE3F-L4010401	<input checked="" type="checkbox"/>	hdisk8	20 GB	U78A5.001.WIH23CF-P1-C11-L1-T1-W500507630E87FE3F-L4010401	<input checked="" type="checkbox"/>	hdisk9	20 GB	U78A5.001.WIH23CF-P1-C11-L1-T1-W500507630E87FE3F-L4010401
Select	Name ^	Size	Physical Location Code																																				
<input type="checkbox"/>	hdisk1	68.37 GB	U78A5.001.WIH23CF-P1-T5-L60000-L0																																				
<input type="checkbox"/>	hdisk2	30 GB	U78A5.001.WIH23CF-P1-C11-L1-T1-W500507630E87FE3F-L4010400																																				
<input type="checkbox"/>	hdisk3	30 GB	U78A5.001.WIH23CF-P1-C11-L1-T1-W500507630E87FE3F-L4010400																																				
<input type="checkbox"/>	hdisk4	30 GB	U78A5.001.WIH23CF-P1-C11-L1-T1-W500507630E87FE3F-L4010400																																				
<input type="checkbox"/>	hdisk6	20 GB	U78A5.001.WIH23CF-P1-C11-L1-T1-W500507630E87FE3F-L4010401																																				
<input type="checkbox"/>	hdisk7	20 GB	U78A5.001.WIH23CF-P1-C11-L1-T1-W500507630E87FE3F-L4010401																																				
<input checked="" type="checkbox"/>	hdisk8	20 GB	U78A5.001.WIH23CF-P1-C11-L1-T1-W500507630E87FE3F-L4010401																																				
<input checked="" type="checkbox"/>	hdisk9	20 GB	U78A5.001.WIH23CF-P1-C11-L1-T1-W500507630E87FE3F-L4010401																																				

Figure 7-27 Disk selection

10. Also depending on the installation preparation in the Optical devices menu, you can select either:
- a. The USB DVD drive to the IBM i 6.1 partition (the media tray is assigned to this blade) in the **Available Physical Optical Devices** section.
 - b. Or a predefined Virtual Optical Device in the **Virtual Optical Devices** section. This shortens the installation time using Virtual Optical Devices with predefined images of the LIC, BOSS_1, and BOSS_2 DVDs installing a second or third IBM i 6.1 partition on the same BladeCenter JS23/JS43 server, and a manual exchange of the install media is omitted.
 - c. Select the first row in the Virtual Optical Devices section.
 - d. Click **Modify** underneath the Current Media field and select the desired Virtual Optical Media.
 - e. Click **Next** to continue.

Figure 7-28 on page 277 provides an example of the optical selection panel. If the media tray for the BladeCenter has been assigned to the blade server you are creating the partition on, the device will be available. Under the *Physical Optical Devices* area is the checkbox to select cd0.

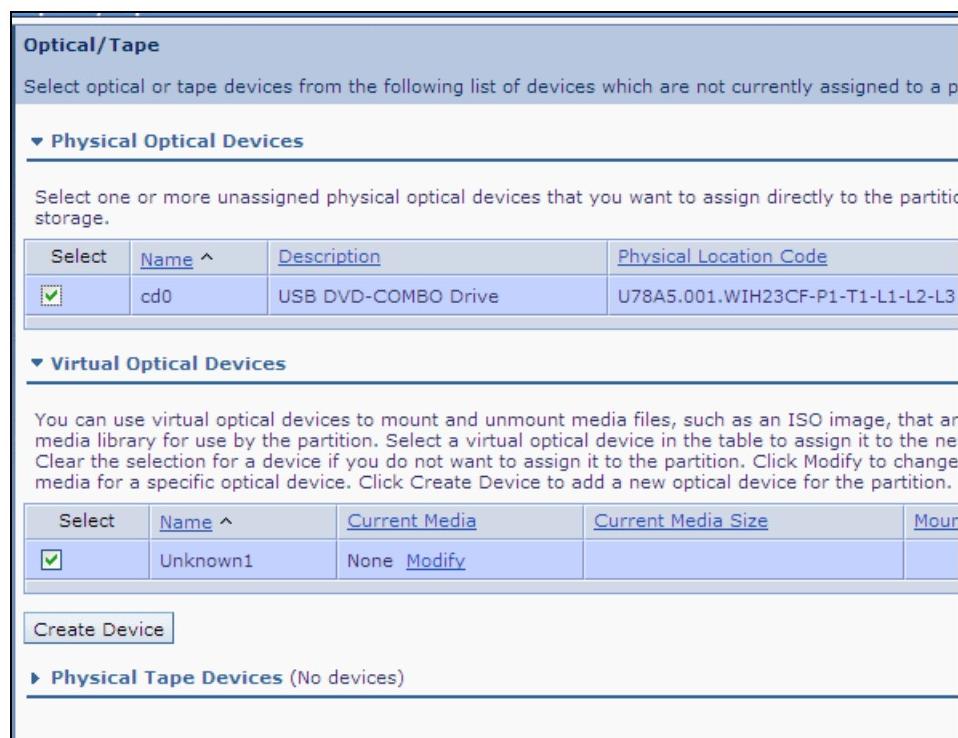


Figure 7-28 Partition optical selections

11. Review the summary of your definition and click **Finish** to create the IBM i 6.1 partition.

7.3.4 Partition properties

Earlier in the process of creating the IBM i partition, selections were made for processor and memory values. With the partition create wizard there are not options to tailor the memory and processor values like there are when using an HMC. From the Memory and Processor tabs you can change the selections to allow a minimum and maximum value. This allows some movement of memory and processor resources using dynamic allocations.

Figure 7-29 on page 278 shows the fields for *Load Source and Console Identification*. This information can be found by selecting the partition and using

the Properties task. The first tab of the properties box is the General tab. Here you can view the fields for the load-source adapter and the console adapter. The selections should be the virtual adapters when in the blade environment. Also note that the IPL source will be set to D which uses the *Alternate restart adapter*.

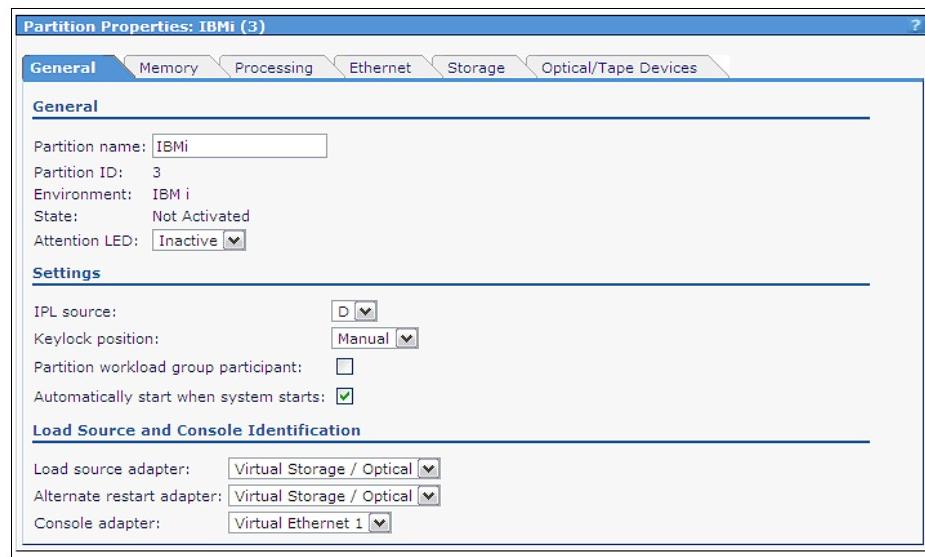


Figure 7-29 Load Source and Console Identification fields

Figure 7-30 on page 279 provides an example of the Memory tab. You can adjust the partition memory allocation by changing the values and clicking OK. By setting the minimum and maximum values you can create a range of memory to stay within when performing dynamic allocation. Changing the *Assigned* value and clicking on OK would dynamically adjust the memory for the partition provided you stay within the minimum and maximum values. You can change the minimum and maximum values as well but it will cost an IPL of the partition to set the new values.

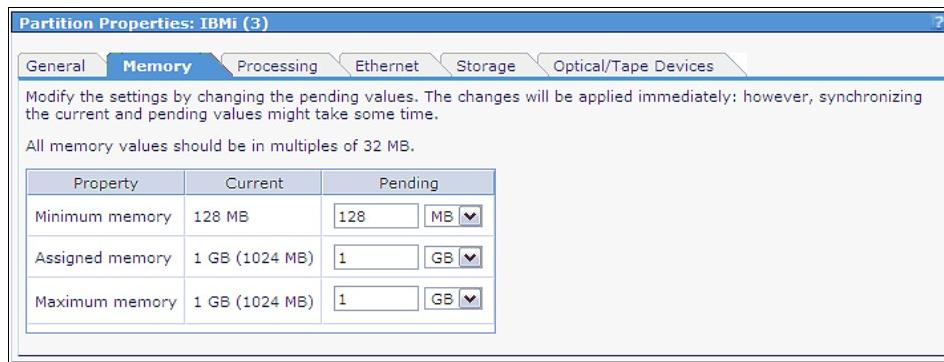


Figure 7-30 Partition memory allocation

Figure 7-31 on page 279 provides an example of the *Processing* tab. You can adjust the partition processor allocations by changing the values and clicking OK. Like the memory, you can adjust the minimum and maximum values to create a range of processor units to stay within when performing dynamic allocation.

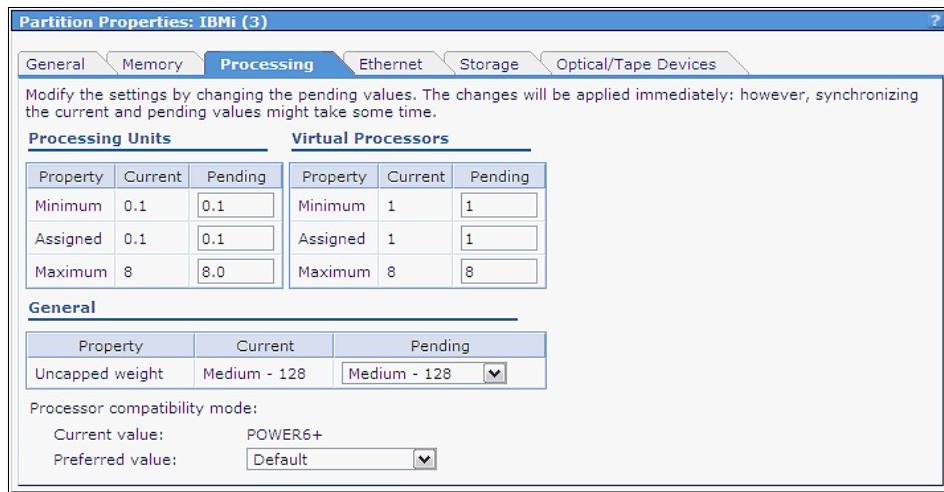


Figure 7-31 Partition processing properties tab

Figure 7-32 on page 280 shows an example of a modified set of values for processing units. In this example the values were modified to have 1.2 Assigned units and a maximum value of 4.0 units. Also notice that *Virtual Processors* values were modified. With these settings we can have anywhere from 1 processing units to 4.0 processing units. Anytime a change is desired within the minimum and maximum values, that change can be performed dynamically by modifying the *Assigned* value.

Processing Units			Virtual Processors		
Property	Current	Pending	Property	Current	Pending
Minimum	0.1	0.1	Minimum	1	1
Assigned	0.1	1.2	Assigned	1	2
Maximum	8	4.0	Maximum	8	4

Figure 7-32 Processing units value change

7.3.5 IBM i 6.1 install media preparation

There are two general methods for installing IBM i Licensed Internal Code (LIC) and the 6.1 operating system on a BladeCenter JS32/JS43 blade in an IBM BladeCenter chassis. You can use the CD/DVD drive in the IBM BladeCenter chassis Media module attached to the IBM i 6.1 partition, or you can create virtual optical media devices.

Using virtual optical media devices is very practical during the installation process of LIC, operating system, or related software because the entire contents of the required CD/DVDs is dumped to a virtual optical image. Therefore, during the installation process you do not need to replace the CD/DVD media physically if a second CD/DVD media is required.

For further information about how to create a virtual optical media device, refer to “Creating Virtual Media Library using IVM” on page 322.

One other option is to use a remote PCs CD/DVD device and configure it to the AMM in the blade center. This option would allow you to remotely install your operating system from your PC. See “Attaching a remote PC file or Media device” on page 327 for more information.

7.3.6 Connect the System i LAN console

Connecting a local console on a network (LAN) to a system enables you to have an active console and a functional remote control panel. As previously mentioned, to establish a connection to the IBM i 6.1 partition in the BladeCenter S chassis, the PC must be in the same TCP/IP subnet as the BladeCenter JS23/JS43 blade.

If you experience problems establishing a connection, check the settings of the installed firewall. The ports used by the System i LAN console may have been disabled by the firewall.

To establish a connection to an IBM i 6.1 created partition, follow these steps:

1. In the active Windows session, select **Start** → **All Programs** → **IBM System i Access for Windows** → **Operations Console**. The window shown in Figure 7-33 on page 281 should appear.

Select the **Connection** drop-down menu to continue.

Figure 7-33 on page 281 shows an example of the System i Operations Console panel.

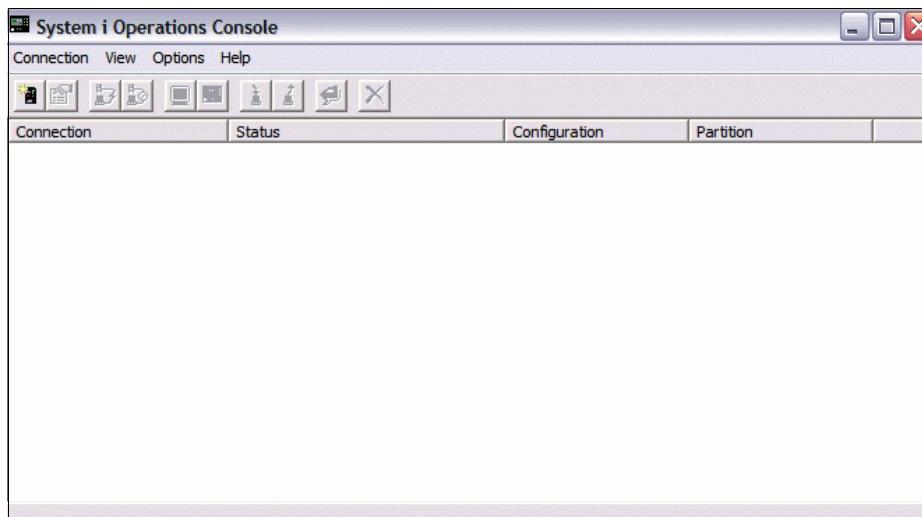


Figure 7-33 IBM System i Operations Console

You will reach the window shown in Figure 7-34 on page 281.

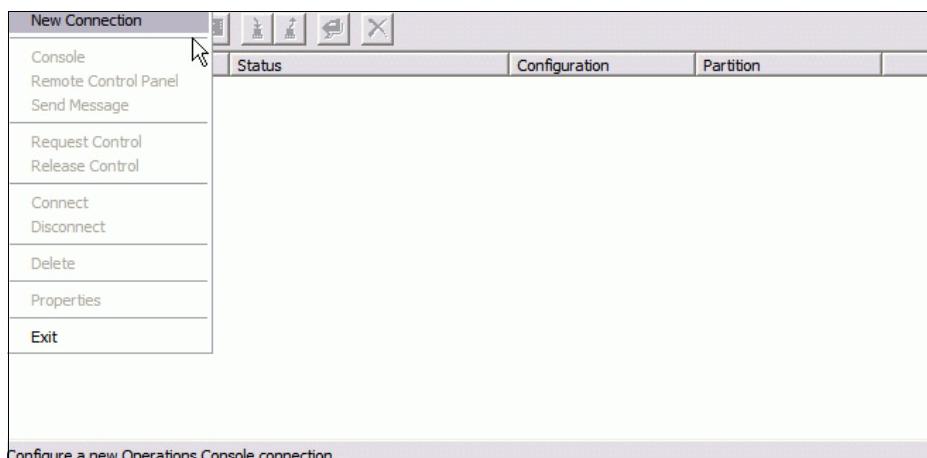


Figure 7-34 IBM System i operator console

2. Click **New Connection** to continue.

You reach the Operation Console Configuration wizard Welcome screen, as shown in Figure 7-35 on page 282.

A connection to the Internet is required to reach the InfoCenter services.

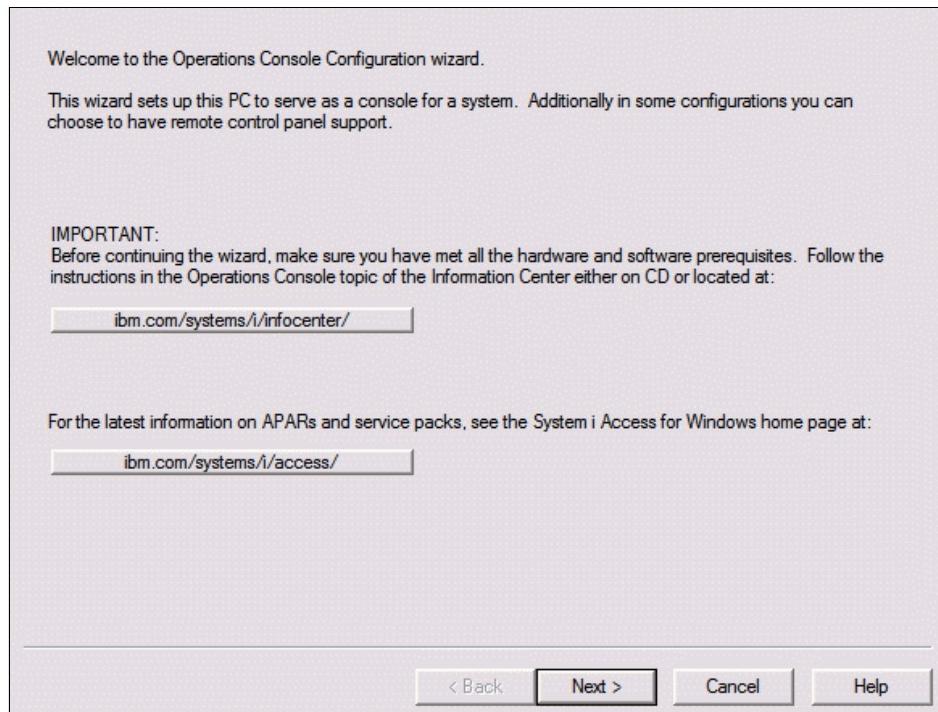


Figure 7-35 IBM System i Operations Console Welcome

Click **Next** to continue.

You may also see a dialog box appear asking you to confirm that the prerequisites for Operations Console have been met. Clicking on the **Help** button will provide the needed information. If the prerequisites have been met then click **Yes** to continue.

3. Depending on the actual infrastructure or existing systems, a connection can be made by using a direct connection or a LAN connection, as shown in Figure 7-36 on page 283.

In our case we selected the method **Local console on a network (LAN)**.

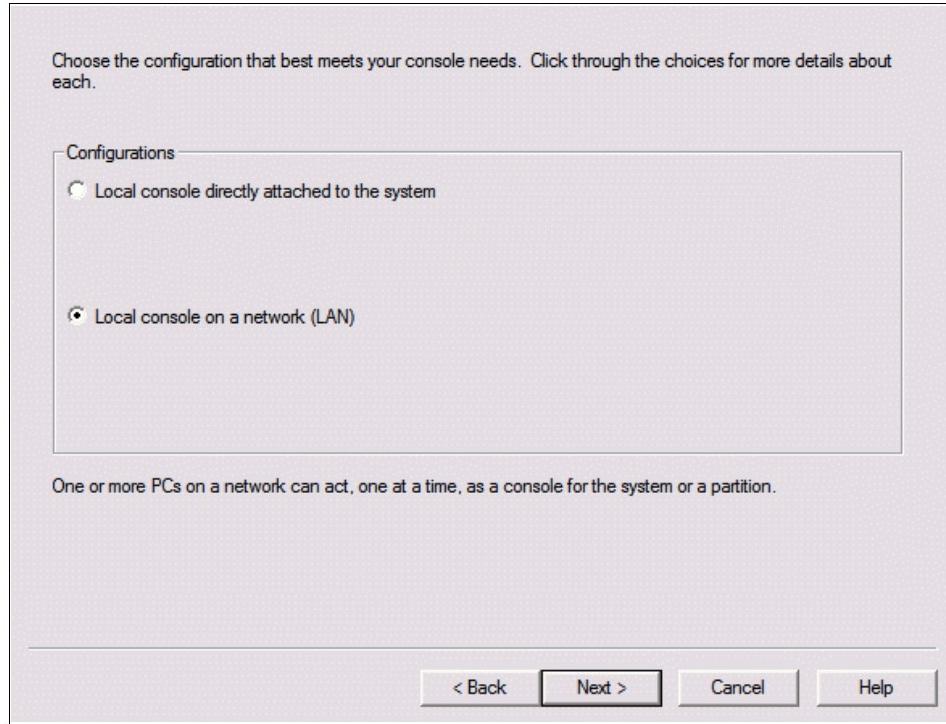


Figure 7-36 IBM System i Operations Console - choose a configuration

Click **Next** to continue.

4. The System i service hostname must be defined first to establish a connection to the BladeCenter JS23/JS43 blade; see Figure 7-37 on page 284.

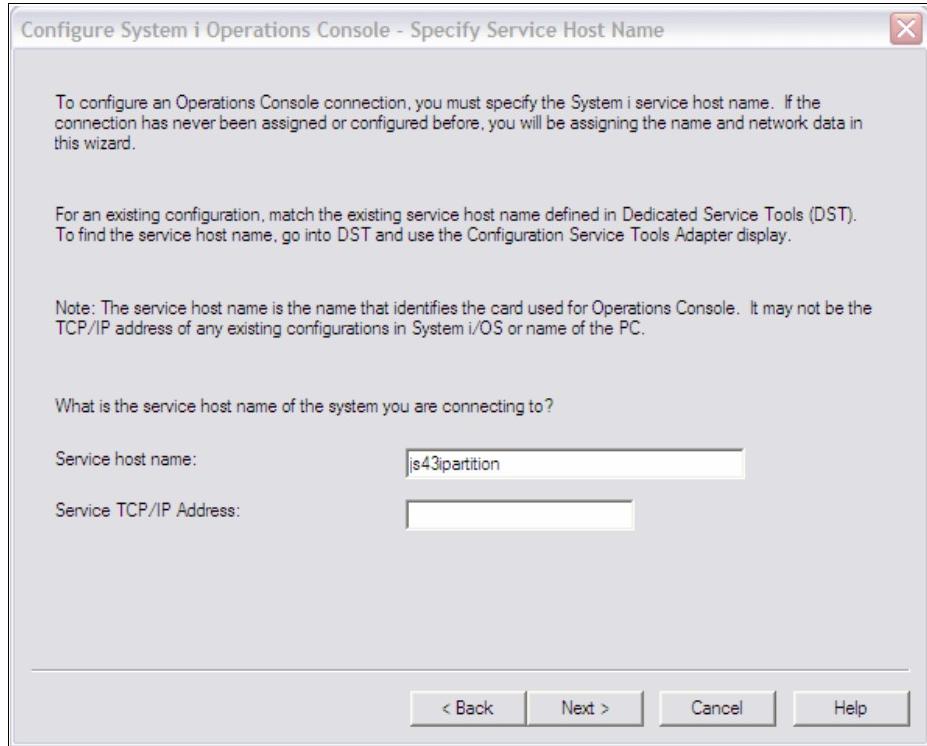


Figure 7-37 IBM System i Operations Console - enter the Service host name

Enter the service host name and click **Next**.

The System i service host name (interface name) is the name that identifies the service connection on your network that is used for service tools, which includes an Operations Console local console on a network (LAN) configuration. This is assigned by your system or network administrator and must be resolved through DNS. Since this is the first connection, the service host name information will be passed to the service tools LAN adapter.

You need a service host name (interface name) any time a console or remote control panel is being connected using a network connection. For example, if a server is logically partitioned, then even though the primary may have a non-networked console, having a remote control panel to a secondary partition may be desirable.

5. If this is a first-time connection, the console will respond as shown in Figure 7-38 on page 285.

If the Service host name is already defined in the host file of the console PC, then the window will resolve the IP address automatically and fill in the field as shown in Figure 7-39 on page 286.

Note: Choose a service host name that is related to the IBM i 6.1 partition name created in Integrated Virtualization Management (IVM) so that you can more easily remember which partition is meant.

The service host name and service TCP/IP address are stored automatically in the host file of the IBM System i Access for Windows console PC.

For Windows XP environments you will find the hosts file at the following path:
C:\WINDOWS\system32\drivers\etc\hosts

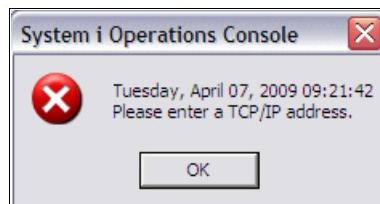


Figure 7-38 IBM System i operator console

Click **OK** to continue.

6. Enter the service TCP/IP address in the field shown in Figure 7-39 on page 286.

Important: Do not use the IP address from the VIO Server here. The System i Access for Windows software establishes a connection to the IBM i partition through the VIOS.

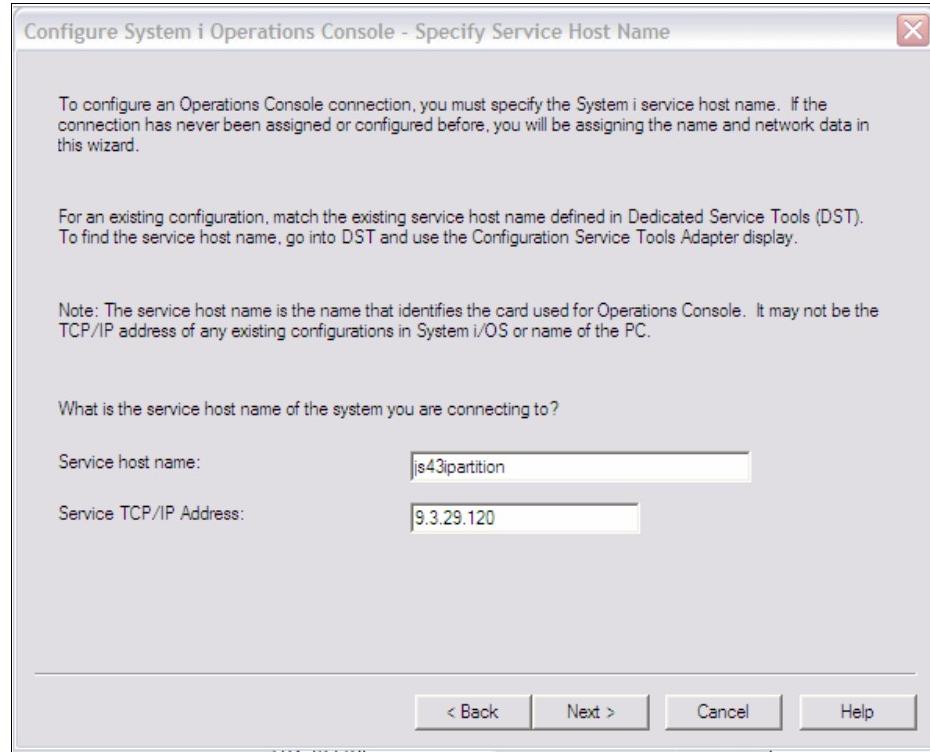


Figure 7-39 IBM System i Operations Console - enter the Service TCP/IP Address

Click **Next** to continue.

7. After providing an IP address, the wizard will now ask for the subnet mask, gateway, serial number and partition id as shown in Figure 7-40 on page 287

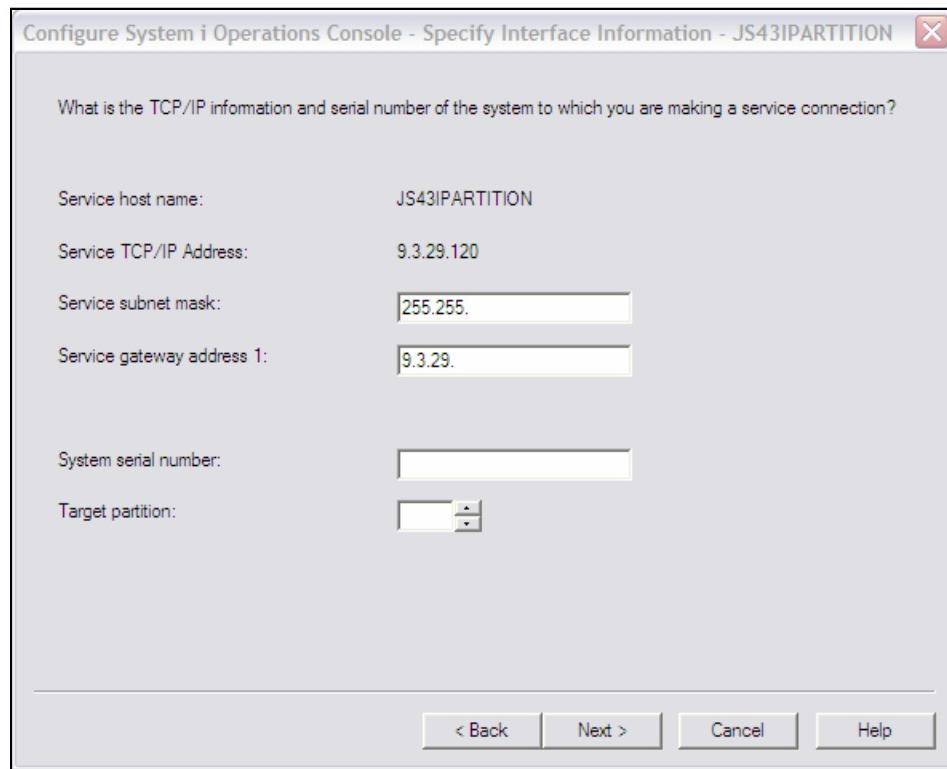


Figure 7-40 Specify Interface Information

8. Modify the required fields to the actual implementation. In our hardware scenario a gateway was implemented. Two important fields are System serial number and Target partition, as shown in Figure 7-43 on page 289.

System serial number This is the BladeCenter JS23/JS43 unique system number. To find the System serial number, use the Integrated Virtualization Management (IVM) console and look under System Properties. An example is shown in Figure 7-41 on page 288.

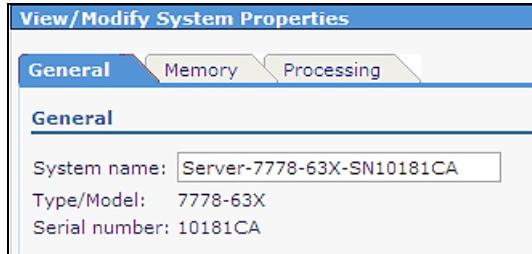


Figure 7-41 System Properties - Serial number

Target partition

This is the Target partition is the partition ID of the IBM i 6.1 partition. To see if partition ID 1 is predefined to VIOS, use IVM. If no other partition is created at this time, the IBM i 6.1 partition ID is 2. The partition ID can be found by looking at the View/Modify partition panel. Next to the partition name is the ID field as shown in Figure 7-42 on page 288. In our example the partition ID is 3.

Partition Details			
Select	ID ^	Name	State
<input type="checkbox"/>	1	is43-vios	Running
<input type="checkbox"/>	2	is43-lp1	Not Activated
<input type="checkbox"/>	3	IBMi	Not Activated
<input type="checkbox"/>	4	testlpar	Open Firmware

Figure 7-42 Partition ID

Enter the appropriate values and click **Next** to continue.

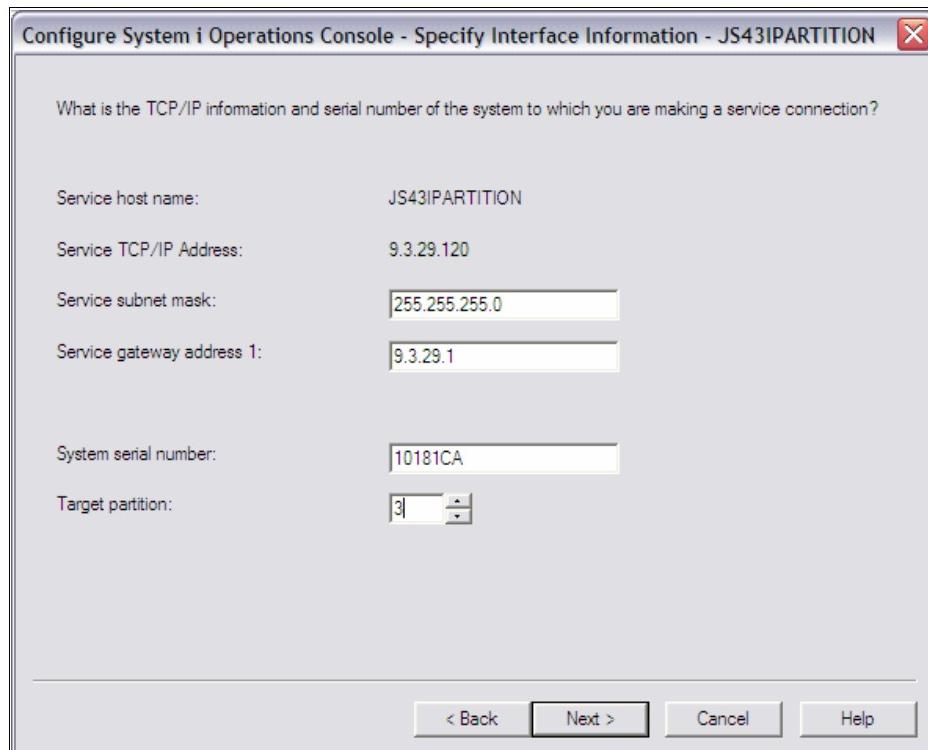


Figure 7-43 IBM System i Operations Console - enter System serial number

Enter values and click **Next** to continue.

9. The next window that appears requests a Service tool device ID to authenticate the communication between the LAN console PC and the IBM i partition, as shown in Figure 7-44 on page 290.

Service tool user IDs are user IDs that are required for accessing service functions through dedicated service tools (DST), system service tools (SST), the System i Navigator (for logical partitions and disk unit management), and the Operations Console. QCONSOLE is the default service tool user ID to be used with console connections. It is however recommended to use ID 11111111 since it is the only service tools ID that is not set as expired.

Service tools user IDs are created through DST or SST and are separate from IBM i 6.1 user profiles. It is also recommended to create additional backup service tool IDs after you have installed your partition. For more information about this topic, refer to Service tools user IDs at:

<http://publib.boulder.ibm.com/infocenter/systems/scope/i5os/topic/rzamh/rzamhwwhatuserids.htm>

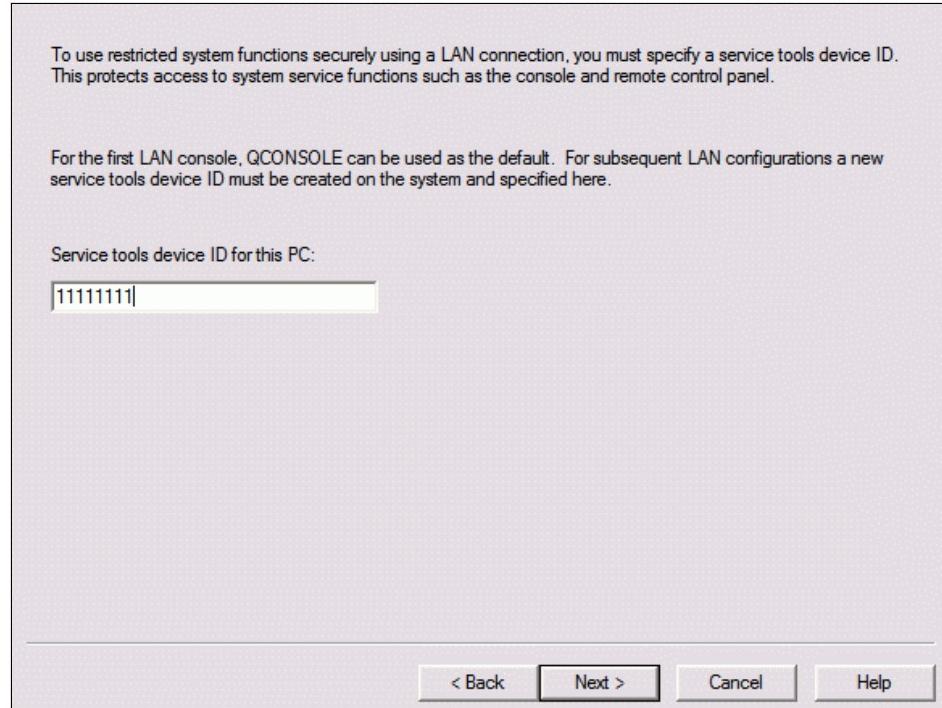


Figure 7-44 IBM System i Operations Console - enter Service tools device ID

Enter the **Service tool device ID** and click **Next** to continue.

10. Figure 7-45 on page 291 shows the final window that is displayed after you define the recommended information for an IBM System i Operations Console.

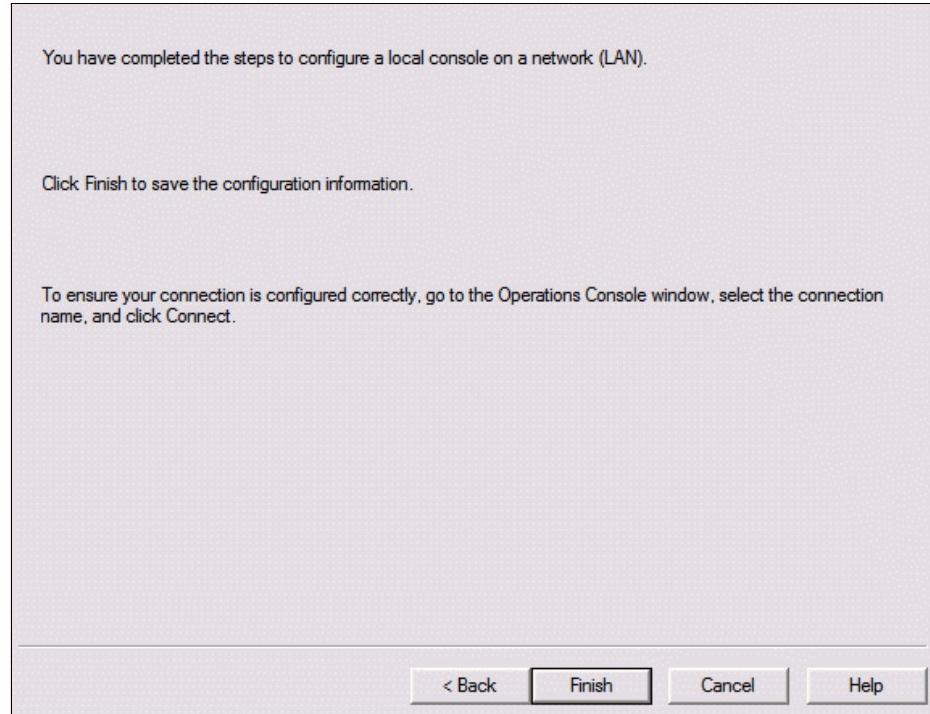


Figure 7-45 IBM System i Operations Console - finalizing the setup

Click **Finish** to save the configuration information.

The configuration window will close immediately and you will return to the initial window with the predefined console definitions for a BladeCenter JS23/JS43 blade, as shown in Figure 7-46 on page 291.

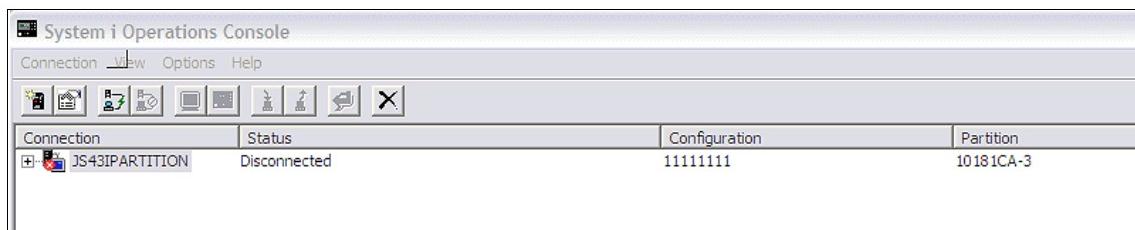


Figure 7-46 IBM System i Operations Console

To connect the IBM System i Operations Console to the IBM i 6.1 partition, click on the connection name, then click on the connect icon or use the **Task Connection → Connect**. Figure 7-47 on page 292 shows an example of the

connect icon. Once the session starts the connection the partition can be activated. Partition activation is discussed in the next section.

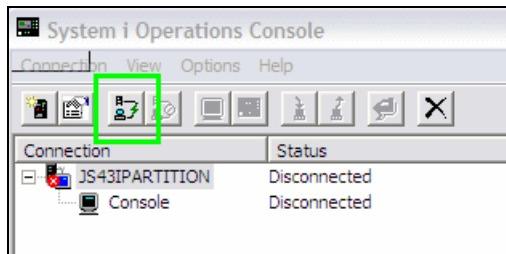


Figure 7-47 Connect console session

7.3.7 IBM i 6.1 IPL types

The IPL type determines which copy of programs your system uses during the initial program load (IPL).

- | | |
|-------------------|---|
| IPL type A | Use IPL type A when directed for special work, such as applying fixes (PTFs) and diagnostic work. |
| IPL type B | Use the B copy of Licensed Internal Code during and after the IPL. This copy resides in System Storage Area B. This copy contains temporarily applied fixes. |
| IPL type C | Development support reserves this type of IPL for hardware service representatives.

Note: Do <i>not</i> use this function. Data loss can occur with improper use of this function. |
| IPL type D | Use IPL type D when directed for special work, such as installing and reloading programs. IPL type D loads the system programs from an alternate IPL load source, such as a tape drive or CD-ROM.

Typically, an IPL uses programs that are stored on the primary IPL load source (typically a disk drive). However, sometimes it is necessary to perform an IPL from another source, such as programs that are stored on tape. To do this, you must use IPL type D to perform an IPL from the alternate IPL load source. |

When configuring the IBM i 6.1 partition, use IPL Type D to install and set up the environment. After the Licensed Internal Code (LIC) is successfully installed, the installation process will automatically change the IPL type to A.

Note: Typically after installation of PTFs you will run the partition on the B side. This value is changed on the General tab of the partition properties.

After the prerequisites are completed, the steps required to install 6.1 on a BladeCenter JS23/JS43 are essentially the same as on any other supported system:

1. Place the IBM i 6.1 installation media in the DVD drive in the BladeCenter media tray, which at this point should be assigned to your BladeCenter JS23/JS43. Or use the previous created Virtual Optical Media device described in 7.5.3, “Adding Image Files to Media Library” on page 324.
 2. In IVM, select the View/Manage partitions task.
 3. Select the IBM i 6.1 partition and click **Activate**.
 4. After connecting the IBM System i for Windows LAN console, the screen shown in Figure 7-48 appears. The default language feature for English is 2924. This can be changed on the next screen.

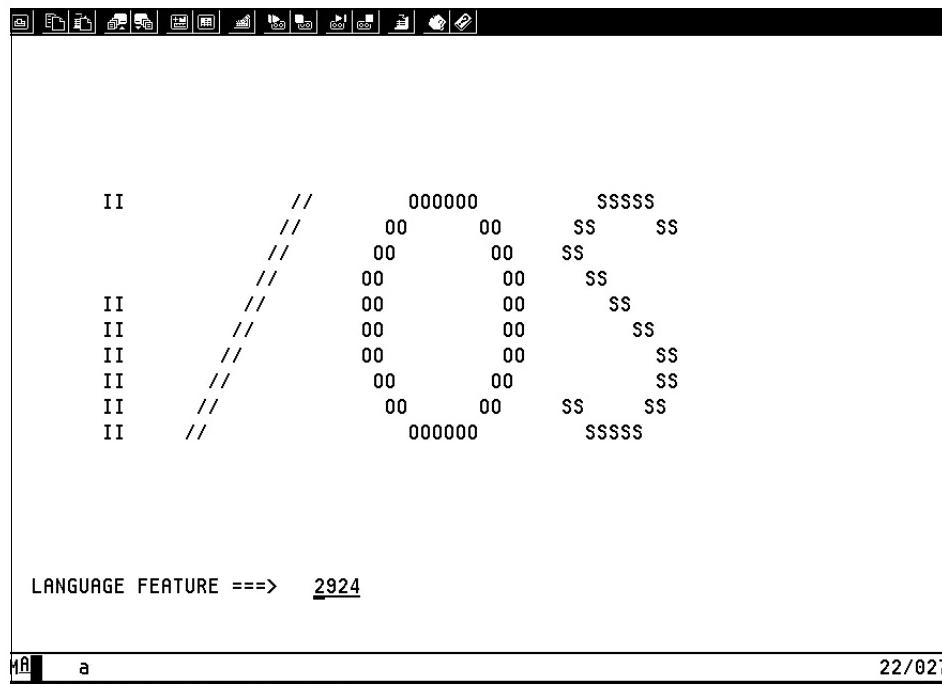


Figure 7-48 LIC initial installation screen

Press **Enter** to continue.

5. Depending on the native language, a selection can be made in the following screen as shown in Figure 7-49 on page 294.

Normally the same language will be chosen as the language for the IBM i 6.1 operating system. Language feature 2924 enables the English environment.

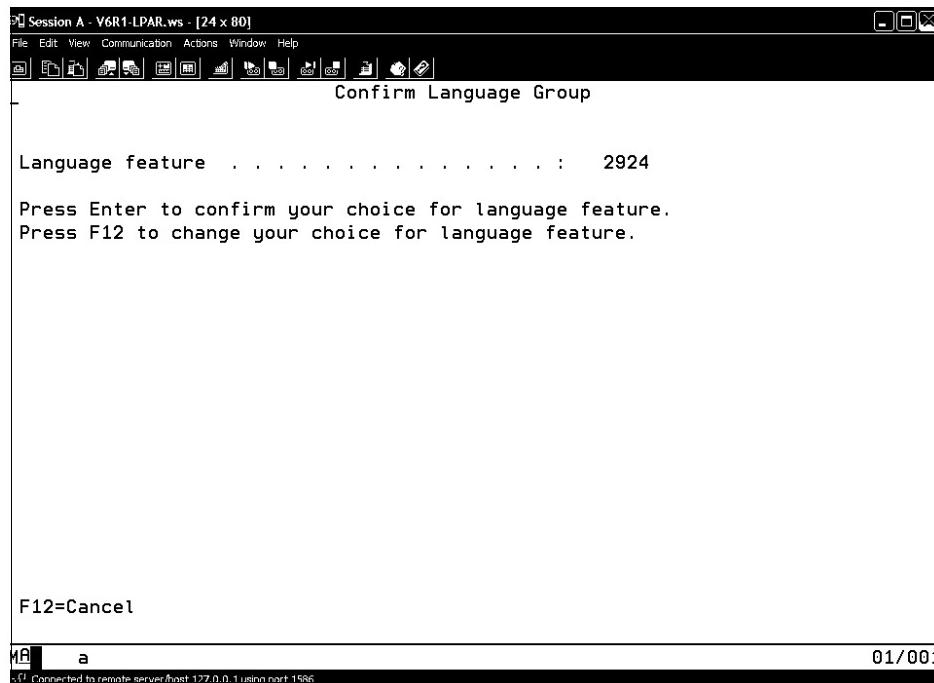


Figure 7-49 Confirm Language setup

Press **Enter** to continue.

The next screen displays several options, as shown in Figure 7-50 on page 295. To install the Licensed Internal Code, type 1 and press Enter.

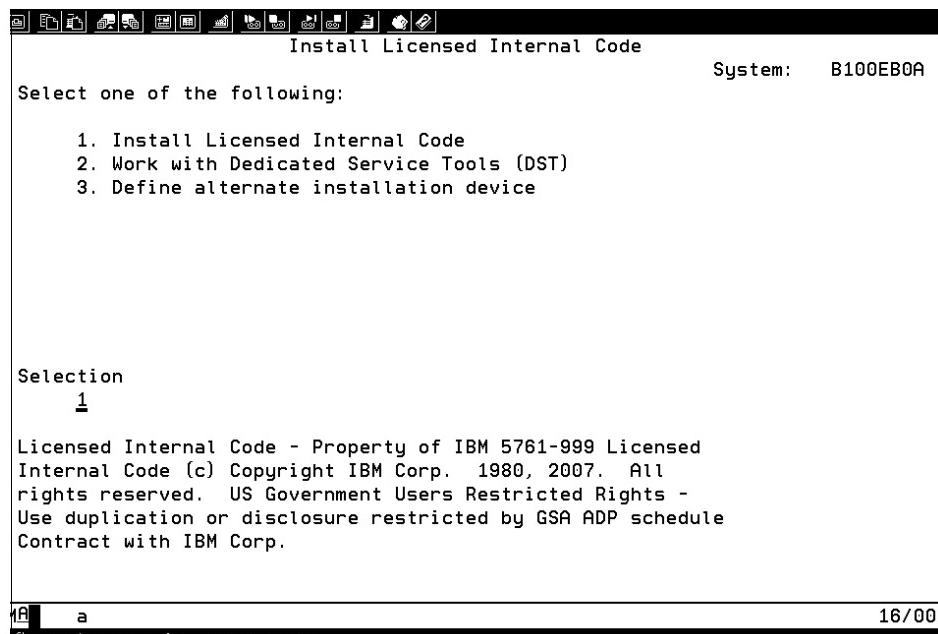


Figure 7-50 *Install LIC*

6. Now select the target install device. Move the cursor to the target device, type 1 and press **Enter**; see Figure 7-51 on page 296.

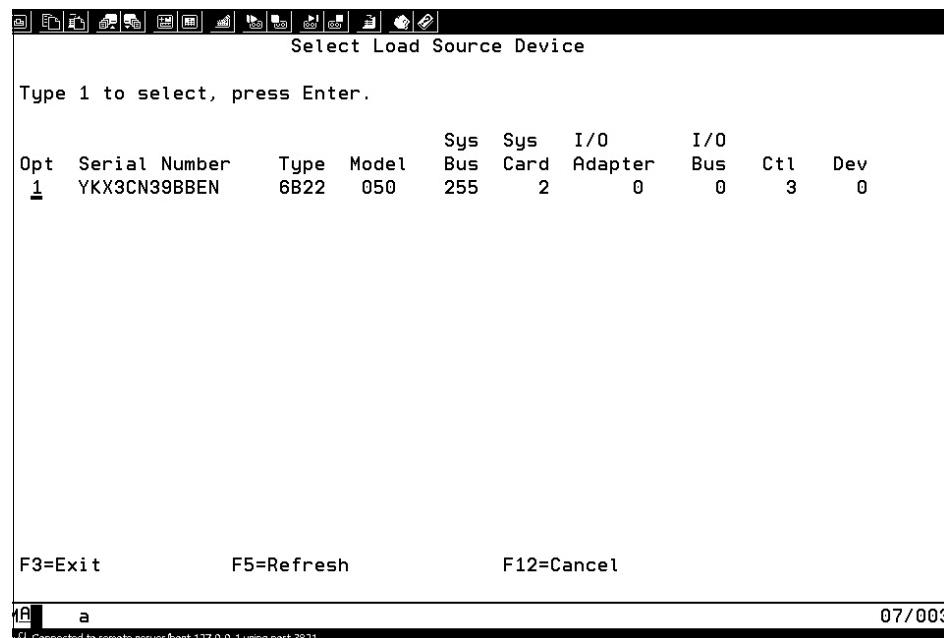


Figure 7-51 Select source disk

7. Confirm the previous selection of the Load Source Device by pressing **F10**; see Figure 7-52 on page 297.

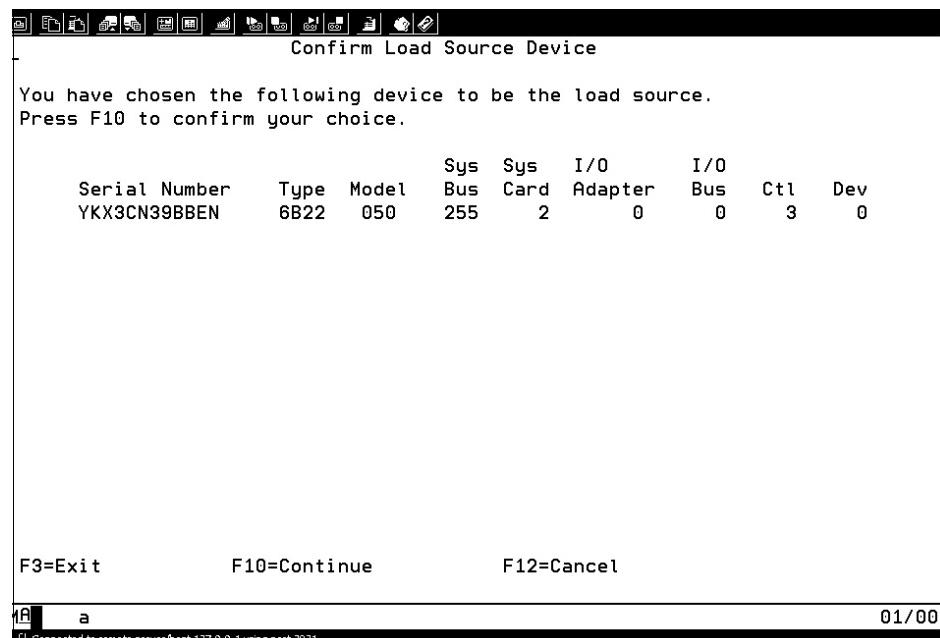


Figure 7-52 Confirm source device

8. The install Licensed Code (LIC) menu appears on the console as shown in Figure 7-53 on page 298. Type 2 for Install Licensed Internal Code and Initialize system, then press **Enter** to continue.

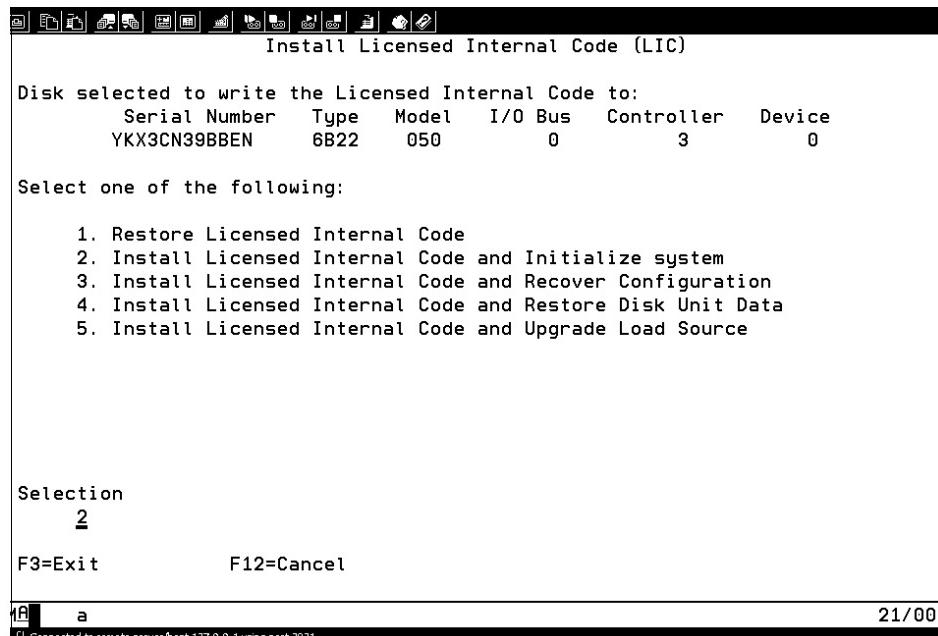


Figure 7-53 Select options

9. The Confirmation screen appears as shown in Figure 7-54 on page 299. This procedure causes existing data on the disk assigned to this logical partition to be lost. Press **F10** to continue or press **F12** to Cancel and return to the previous screen.

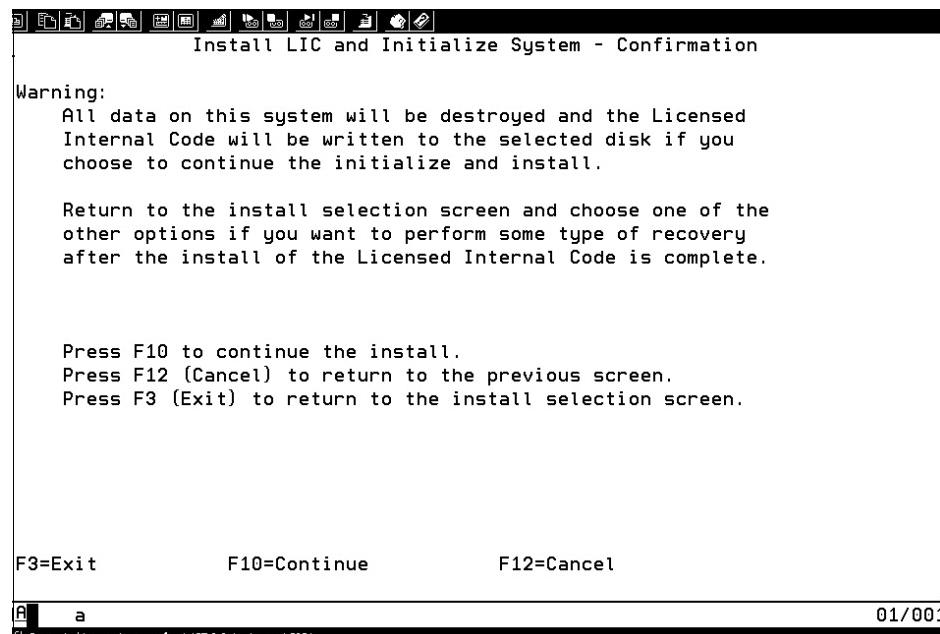


Figure 7-54 Confirm definition

After you confirm the definition, you reach the Initialize the Disk status screen as shown in Figure 7-55 on page 300. Depending on the predefined size of the virtual disk, this procedure can take 60 minutes or more.

```
Initialize the Disk - Status

The load source disk is being initialized.

Estimated time to initialize in minutes : 180
Elapsed time in minutes . . . . . : 1.5

Please wait.

Wait for next display or press F16 for DST main menu
01/00
```

Figure 7-55 Initialize disk

10. Next, the Install Licensed Internal Code status display appears on the console as shown in Figure 7-56 on page 301. It will remain on the console for approximately 30 minutes.

Once the LIC has completed installing, the logical partition is automatically restarted to IPL to DST at this time to complete the Licensed Internal Code installation.

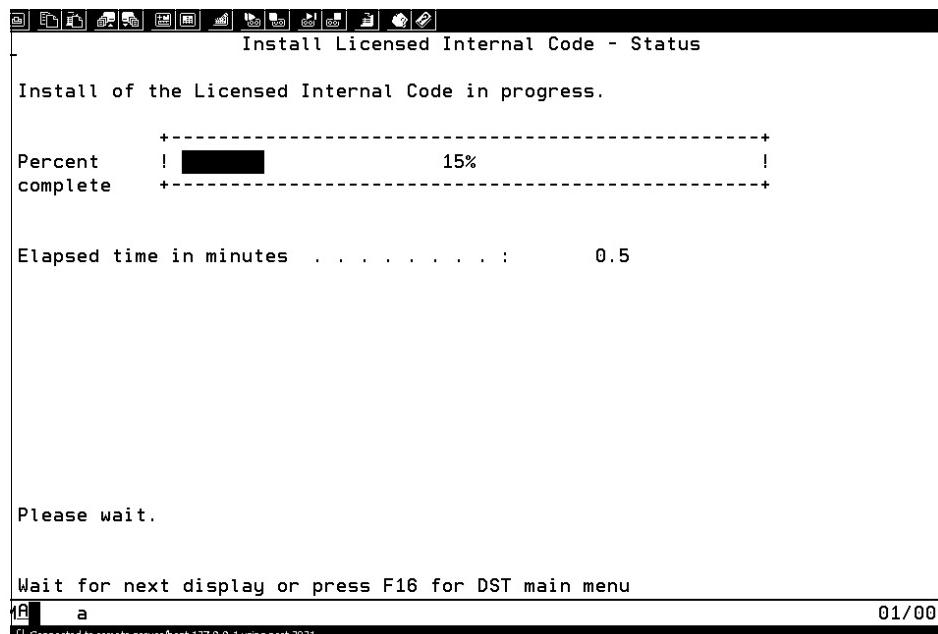


Figure 7-56 Install LIC status

11. The Disk Configuration Attention Report display might appear on the console.

Figure 7-57 on page 302 shows the report for a new disc configuration.

Press **F10** to accept the action to define a new disk configuration.

Note: If the Disk Unit Not Formatted For Optimal Performance Attention Report appears on the console, then further actions should be performed as described in InfoCenter:

<http://publib.boulder.ibm.com/infocenter/iseries/v5r4/index.jsp?topic=/rzahc/rzahcdiskw.htm>

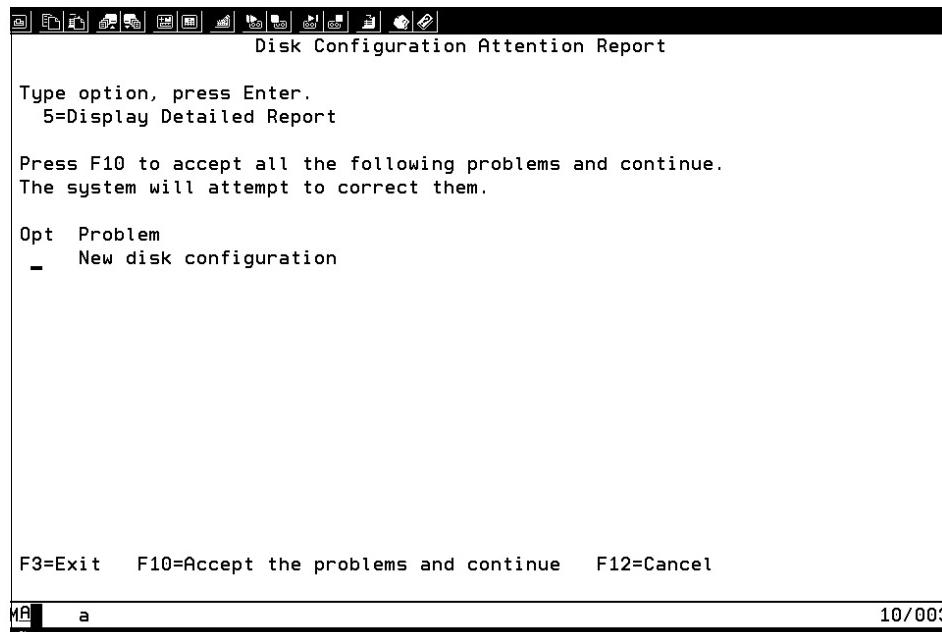


Figure 7-57 Attention Report

After the Licensed internal Code installation is complete, you will see the screen shown in Figure 7-58 on page 303.

At this time it is recommended to complete disk unit configuration before installing the operating system. When completing disk configuration you will be adding additional units and possibly starting mirroring on the disk units.

See the following link to assist with performing disk configuration. Not all steps will need to be performed.

<http://publib.boulder.ibm.com/infocenter/systems/scope/i5os/index.jsp?topic=/rzarm/rzarmrcvaft.htm&tocNode=toc:rzahg:i5os/17/0/5/5/4/>

Note: Disk configuration is not a required action at this time but is recommended in the case of a failure that may cause a reload. It is much quicker to install LIC than it would be to install LIC and the operating system again.

After completing disk configuration you may continue the installation process for IBM i. Select option **2** to install the IBM i operating system after disk configuration steps are completed.

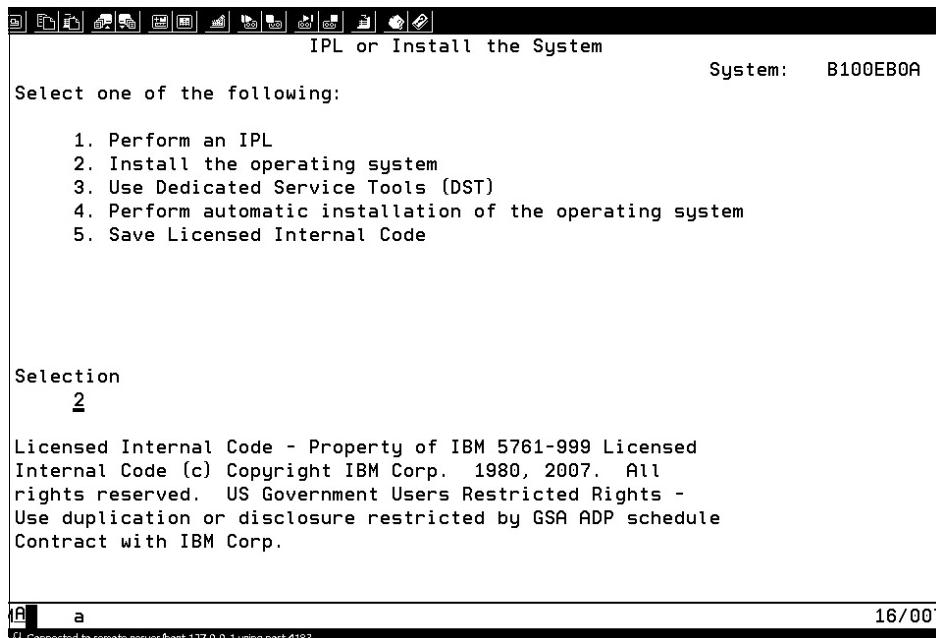


Figure 7-58 Install the operating system

7.4 Installing the IBM i 6.1 Operating System

From the IPL or Install the System screen, the installation process of the operating system can be continued without an interruption. If you use the virtual optical device method of having the two IBM i 6.1 DVDs previously unloaded to virtual optical devices, the only action necessary is to assign the virtual optical device with the IBM i DVD 1 content to the IBM i partition.

1. On the IPL or Install the System screen on the console, type the number that corresponds to the correct device type as shown in Figure 7-59 on page 304. In our case we used option 2 for the assigned virtual optical devices.

If a virtual optical device is not defined, then replace the Licensed Internal Code DVD with the IBM i DVD 1 in the CD/DVD drive in the Media Tray.

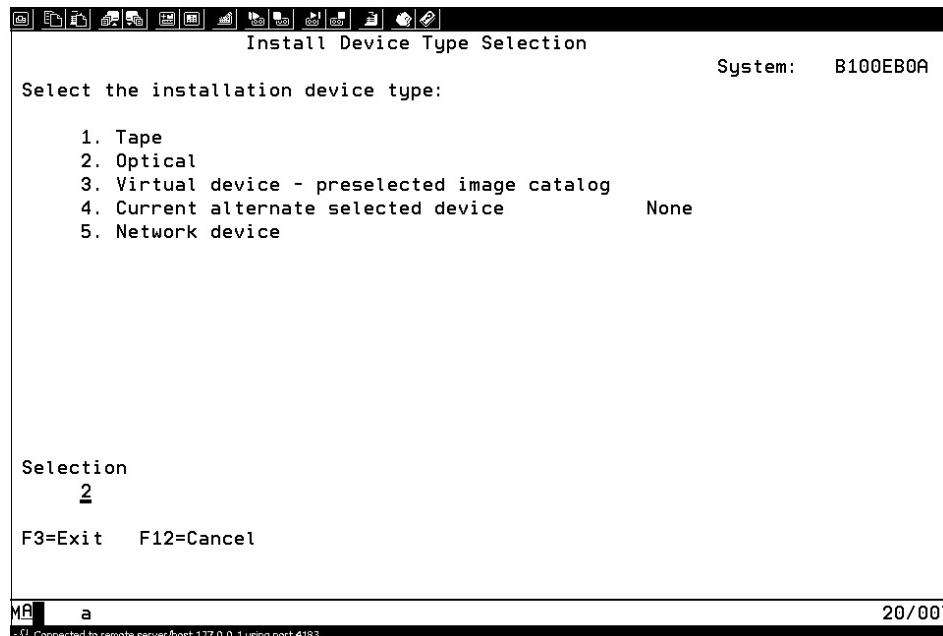


Figure 7-59 Select install device

- Type 2 and press **Enter** to continue.
2. The Confirm Install of the Operating System screen is displayed on the console screen, as shown in Figure 7-60 on page 305.

Press **Enter** to continue the installation process.

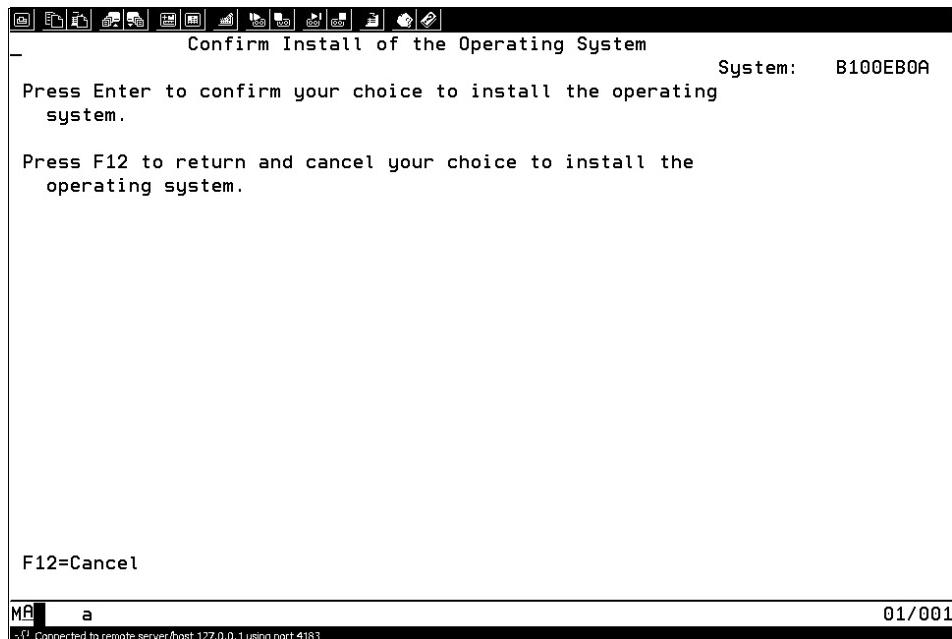


Figure 7-60 Confirm installation

3. The Select a Language Group screen displays the primary language preselection, as shown in Figure 7-61 on page 306. This value should match the Language feature number that is printed in the installation media.

The following URL provides the Language feature codes:

<http://publib.boulder.ibm.com/infocenter/systems/scope/i5os/index.jsp?topic=/rzahc/rzahcnlvfeaturecodes.htm&resultof=%22language%22%20%22features%22%20&searchQuery=language%20features&searchRank=0&pageDepth=0>

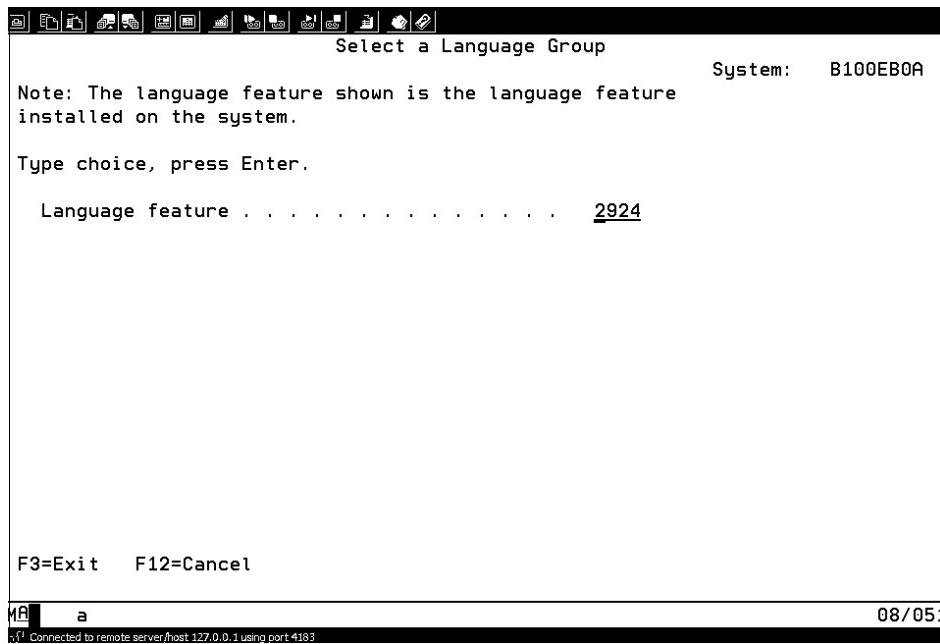


Figure 7-61 Select language feature

4. Type choice and press **Enter** to continue.

The Confirm Language Feature Selection appears on the console, as shown in Figure 7-62 on page 307.

Press **Enter** to confirm and continue.

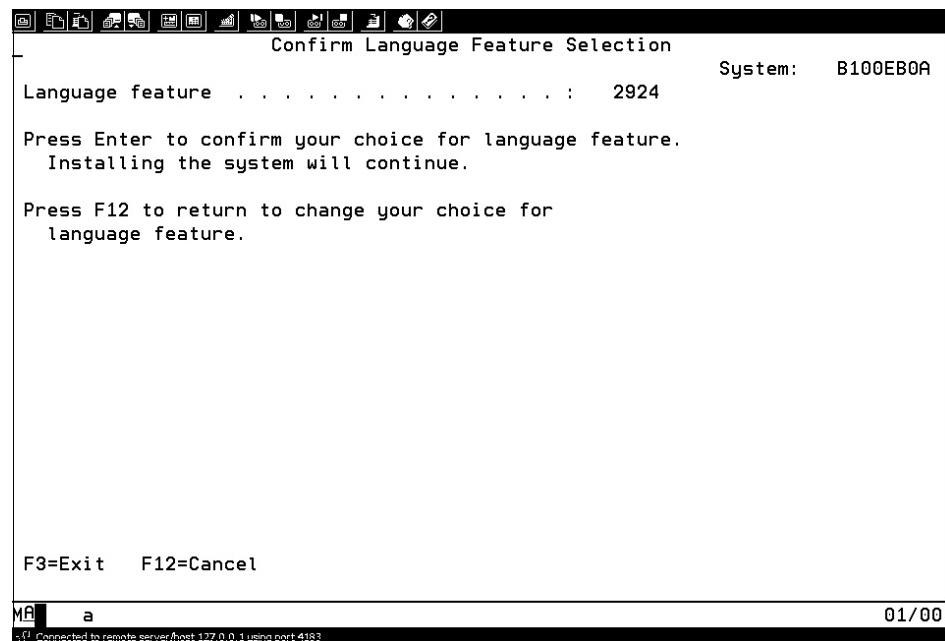


Figure 7-62 LIC install confirm language

5. The Licensed Internal Code IPL in Progress screen appears on the console, as shown in Figure 7-63 on page 308. No administrator action is required.

Licensed Internal Code IPL in Progress
06/17/08 00:03:57

IPL:

Type : Attended
Start date and time : 06/17/08 00:03:23
Previous system end : Abnormal
Current step / total : 12 16
Reference code detail : C6004059

IPL step Time Elapsed Time Remaining
Data Base Recovery 00:00:13 00:00:00
Journal Synchronization 00:00:01 00:00:00
>Commit Recovery
Data Base Initialization
Journal IPL Clean up

Item:
Current / Total :

Sub Item:
Identifier :
Current / Total :

Figure 7-63 IPL in progress

The Install the Operating System screen appears on the console, as shown in Figure 7-64 on page 309.

6. Change the date and time values to the appropriate settings. You must use the 24-hour clock format to set the current time.

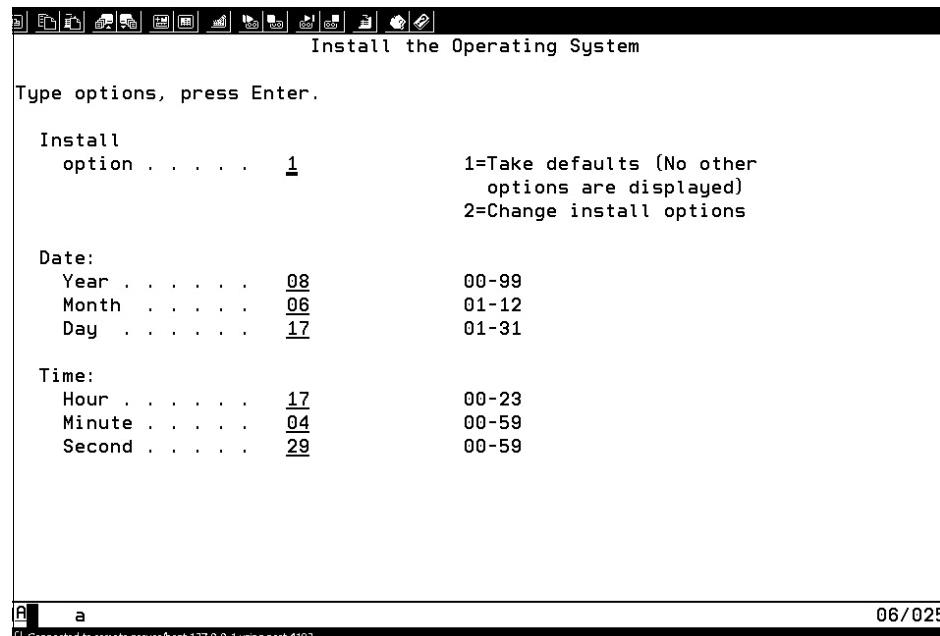


Figure 7-64 Set date and time

7. Figure 7-65 on page 310 shows an example of a status display in the operator console during the installation process. No further action required. Note that the display will be blank for a while between Installation Stage 4 and 5.

The screenshot shows a terminal window titled "Session A - LIVIO.ws - [24 x 80]". The window displays the "i5/OS Installation Status" for Message ID CPI2070. It shows the progress of Stage 2 at 17%, with a bar chart indicating the progress. Below this, it lists the stages of the installation process:

Stage	Completed	Objects Restored
1 Creating needed profiles and libraries :	X	
>> 2 Restoring programs to library QSYS		02272
3 Restoring language objects to library QSYS . . .		
4 Updating program table		
5 Installing database files		-
6 Installing base directory objects.		

At the bottom, it states "i5/OS is a trademark of IBM in the United States and other countries." The bottom status bar shows "MA a X" and "Connected to remote server host 127.0.0.1 using port 1734". The date "17/068" is also visible.

Figure 7-65 Installation status

8. When the Sign On screen is displayed, as shown in Figure 7-66 on page 311, the base installation of the IBM i 6.1 Operating System is finished.

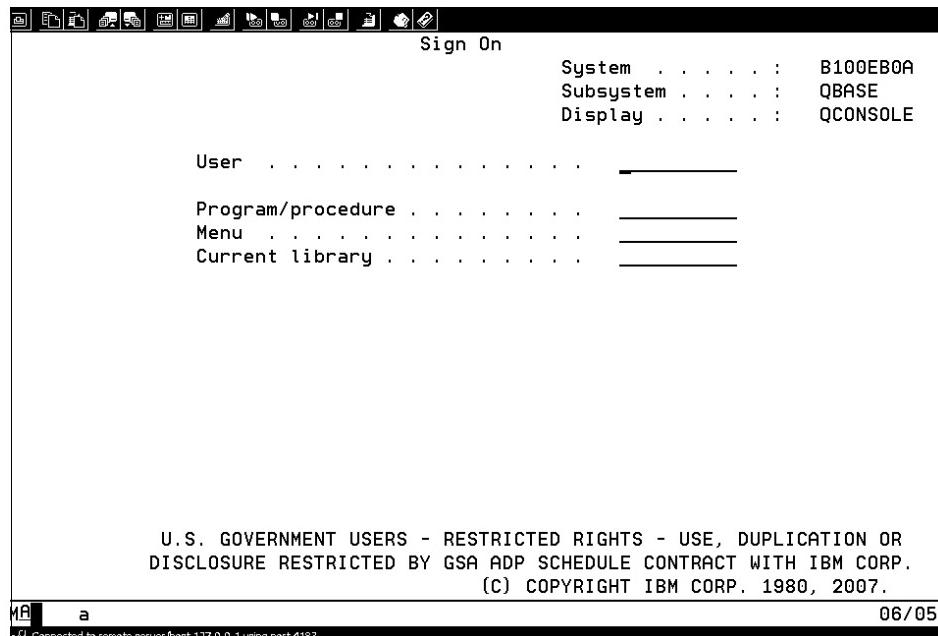


Figure 7-66 Sign On screen

At this stage, the IBM i 6.1 system is ready to use. Information about installing libraries or Licensed Program Products and system configuration is beyond the scope of this book. For detailed software installation information, refer to the following Web site:

<http://publib.boulder.ibm.com/infocenter/systems/scope/i5os/index.jsp?topic=/rzam8/rzam81.htm>

7.4.1 IBM i 6.1 installing PTFs

For detailed explanations about the following tasks, refer to the corresponding sources on the Web.

Use the IBM i recommended Fixes Web site to obtain a list of the latest recommended PTFs:

http://www-912.ibm.com/s_dir/s1kbase.nsf/recommendedfixes

Refer to the primary Web site for downloading fixes for all operating systems and applications:

<http://www-933.ibm.com/support/fixcentral/>

For IBM i 6.1, IBM i5/OS, or OS/400® Operating Systems, fixes are available. To Obtain an IBM i 6.1 fix overview for downloading:

- ▶ Select **System i** in the Product family field.
- ▶ Select **IBM i, i5/OS, and OS/400** in the Product field.
- ▶ Select one the following options in the Ordering option field:
 - Groups, Hyper, Cumulative fixes
 - Individual fixes
 - Search for fixes
- ▶ Select, for example, **V6R1** in the OS level field for fixes for the actual IBM i Operating System version.

Note: To download fixes or obtain information about fixes from the Fix Central Web site, you need a valid IBM ID and password.

7.4.2 IBM i 6.1 TCP/IP setup

Ethernet on a BladeCenter JS23/JS43 server supports TCP/IP, Advanced Peer-to-Peer Networking® (APPN), Advanced Program-to Program Communication (APPC), and remote workstation.

You can use this information to configure and manage TCP/IP on the BladeCenter JS23/JS43 server. Refer to the Ethernet on System i 6.1 topic in the IBM Information Center to configure the IBM i 6.1 communication:

<http://publib.boulder.ibm.com/infocenter/systems/scope/i5os/index.jsp?topic=/rzai2/rzai2configuring.htm&tocNode=toc:rzahg/i5os/10/5/5>

Note: IBM i 6.1 installed on a IBM BladeCenter JS23/JS43 server will communicate to the external LAN using Virtual Ethernet Adapter only, and only those with the Virtual Ethernet bridge enabled.

7.4.3 IBM i 6.1 Navigator for i

The Navigator for i or IBM Systems Director Navigator for i is a Web-based console that consolidates all Navigator for i functions available on the Web. IBM i Systems Director Navigator for i is the tool for IBM i 6.1 management and will work with IBM Systems Director products that enable heterogeneous and cross-system management.

The IBM Systems Navigator for i provides a graphical interface to manage a BladeCenter JS23/JS43 server or Power Systems, as shown in Figure 7-67.

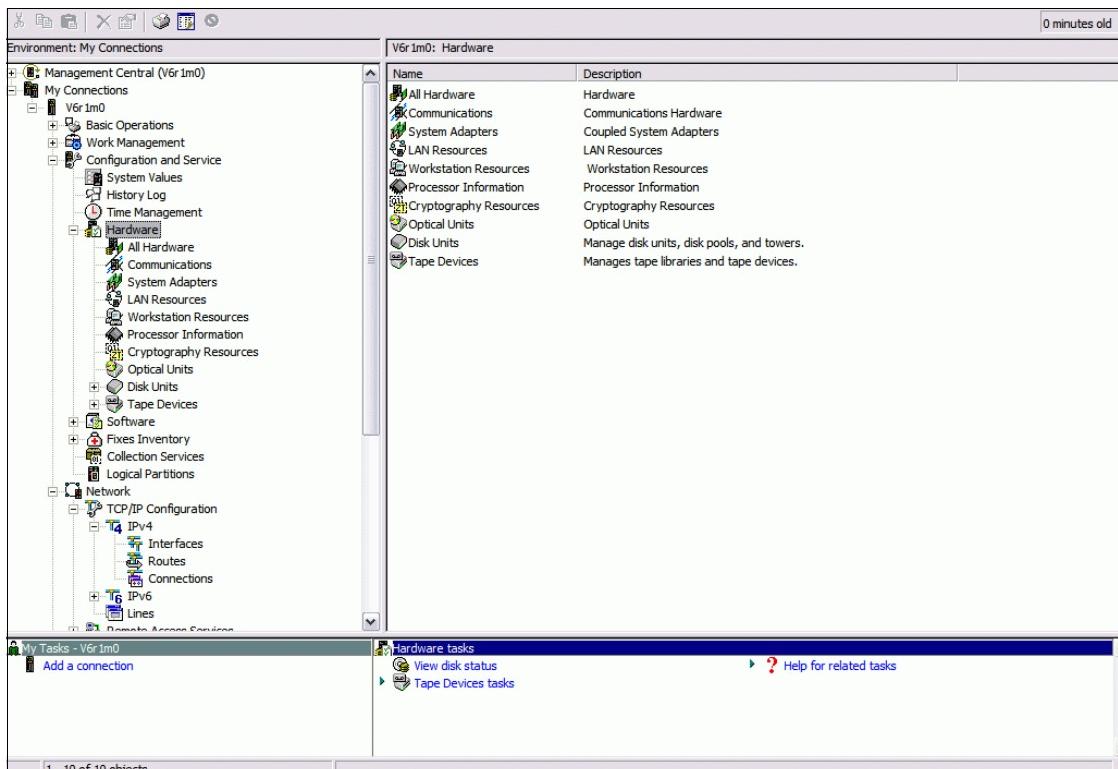


Figure 7-67 IBM Systems Navigator for i

More detailed information to the IBM Systems Director Navigator for i functionality can be found at:

<http://www.ibm.com/systems/i/software/navigator/index.html>

or in *Managing OS/400 with Operations Navigator V5R1 Volume 1: Overview and More*, SG24-6226.

7.5 IBM i 6.1 Backup/Restore

There are two different methods to perform a backup or restore of an IBM i partition.

Important: The virtualized DVD-ROM drive in the chassis *cannot* be used for IBM i 6.1 backups, because it is not writable.

One method is to use file-backed space provided as a virtual optical device. Once the file has been created it can be written to any BSH or BCS supported SAS tape device.

Another method is to use a virtual tape device backed by a SAS tape drive that is virtualized by VIOS. The virtual tape drive method is new as of this writing. This option is only supported using the TS2240 LTO 4 SAS tape drive attached to a SAS switch module. Any other tape drive will not work. When the tape drive is virtualized to the IBM i partition a second virtual SCSI adapter is created for the tape drive connection.

Note: As of this writing, the TS2240 LTO 4 SAS tape drive will be supported for IBM i in the BCH and BCS. This tape drive can be virtualized to the IBM i partition and will appear as a 3580 Model 004.

For further information about using the IBM SAS LTO tape library attached to a BladeCenter S chassis, refer to:

http://www-03.ibm.com/systems/resources/systems_power_hardware_ibm_i_on_b1ade_readme.pdf

IBM i 6.1 backup - virtual optical device

Performing an IBM i 6.1 Operating System backup is a 2-stage process on the IBM BladeCenter JS23/JS43 blade:

1. Create a virtual media library. See section 7.5.1, “Creating a virtual media library for backup” on page 316 for more information.
2. A standard 6.1 **save** command or BRMS (Backup Recovery & Media Services) is used to perform a save on a writable optical device that contains an optical volume. The writable optical device is a file-backed virtual optical drive created in VIOS. The optical volume is a virtual media image, which is a file on disk in VIOS. The file location is `/var/vio/VMLibrary/filename`
3. The virtual media image, containing the 6.1 save, is written out to a SAS attached tape drive using the VIOS command **backup**.

IBM i 6.1 restore - virtual optical device

Performing a restore follows the same 2-stage process in reverse:

1. The virtual media image file is restored from the SAS tape drive onto VIOS disk using the VIOS command **restore**. The image file is then mounted on the correct virtual optical drive assigned to the IBM i 6.1 partition and becomes available as a volume from which to restore.
2. A standard IBM i 6.1 restore is performed from the volume using a **restore** command or BRMS. A full system restore can be executed by first performing a D-mode IPL from the virtual optical image, provided the image contains a full system save performed previously.

IBM i 6.1 backup - virtual tape device

Performing an IBM i 6.1 Operating System backup to virtual tape is a one stage process on the IBM BladeCenter JS23/JS43 blade:

1. Ensure the virtual tape device is assigned to the partition you are performing the backup on. To change or view the assignment use the **View/Modify Virtual Storage** task, then select the **Optical/Tape** tab. Under the **Physical Tape Devices** section select the tape drive and use the **Modify Partition Assignment** button.

Figure 7-68 on page 315 shows an example of the Physical Tape Devices option.

The screenshot shows the Integrated Virtualization Manager (IVM) interface. On the left, there's a navigation sidebar with various management sections like Partition Management, I/O Adapter Management, Virtual Storage Management, IVM Management, System Plan Management, Service Management, and more. The main area is titled 'View/Modify Virtual Storage' and has tabs for Virtual Disks, Storage Pools, Physical Volumes, and Optical/Tape. The Optical/Tape tab is selected. Under 'Physical Optical Devices (No devices)', there's a note about assigning physical optical devices to partitions. Below that is a table for 'Modify partition assignment'. Under 'Virtual Optical Media (No media library)', there's a note about creating a library for virtual optical media. The bottom section, 'Physical Tape Devices', is highlighted with a green border. It has a note about assigning physical tape devices to partitions. Below that is another 'Modify partition assignment' table. This table lists a single entry: 'rmt0' under 'Name', 'Other SAS Tape Drive' under 'Description', 'IBM i (2)' under 'Assigned Partition', and 'U78A5.001.WIH337B-P1-T5-LC000-L0' under 'Physical Location Code'. There are also 'Select' and 'Create Library' buttons.

Figure 7-68 Physical Tape Devices

2. A standard 6.1 **save** command or BRMS is used to perform a save on the tape device (tap0x). If autocfg is on the tape device will configure as an 3580 model 004.

IBM i 6.1 restore - virtual tape device

Performing a restore follows the same 2-stage process.

1. Ensure the virtual tape device is assigned to the parathion you are performing the backup on. To change or view the assignment use the **View/Modify Virtual Storage** task, then select the **Optical/tape** tab. Under the **Physical Tape Devices** section select the tape drive and use the **Modify Partition Assignment** button.
2. Use the standard 6.1 restore command and restore the image from the tape device. The tape device can also be used to perform a D-IPL and full system restore.

7.5.1 Creating a virtual media library for backup

This section describes how to create a virtual media library for storing the IBM i 6.1 operating system backup.

1. Logon the Integrated Virtualization Manager on the BladeCenter JS23/JS43.
2. Select **View/Modify Virtual Storage**, the **Storage Pool** tab, and then click **Create Storage Pool**. The window displayed in Figure 7-69 is shown.

Note: It is recommend to use a separate storage pool. However, rootvg can be used if there is enough free space.

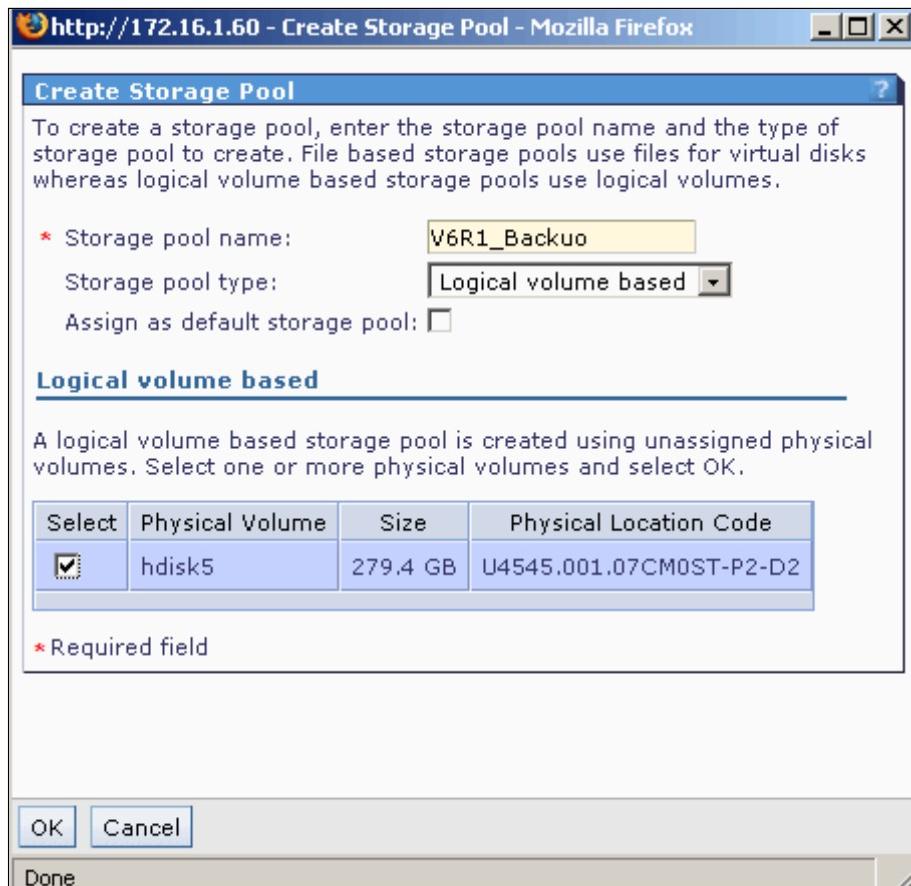


Figure 7-69 IVM Create Storage Pool

3. Enter a name for the storage pool (in our case the internal disk in the BladeCenter S disk module was used), or in a SAN environment, a predefined LUN. Click **OK** to continue.
4. To create the virtual media library click the **Optical Devices** Tab and select **Create Library**.
5. Select the name of the new storage pool and enter an appropriate size for the media library. Select **OK** to continue.
6. To add a new virtual optical device to the media library, select **Add Media** in the Virtual Optical Media section. The window displayed in Figure 7-70 will appear.

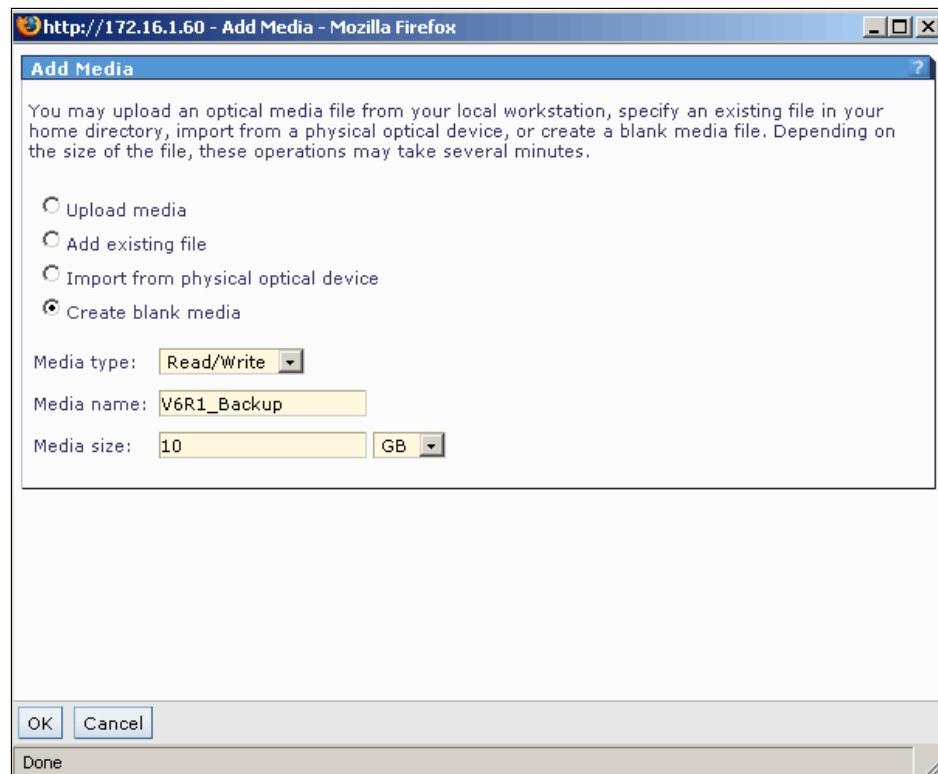


Figure 7-70 IVM Create blank media

7. Select **Create blank media** and enter a meaningful Media Device name and an appropriate size for the new volume. Ensure the media type is set for read/write. Click **OK** to continue.
8. The new virtual optical device should be listed in the Virtual Optical device list, as shown in Figure 7-71 on page 319.

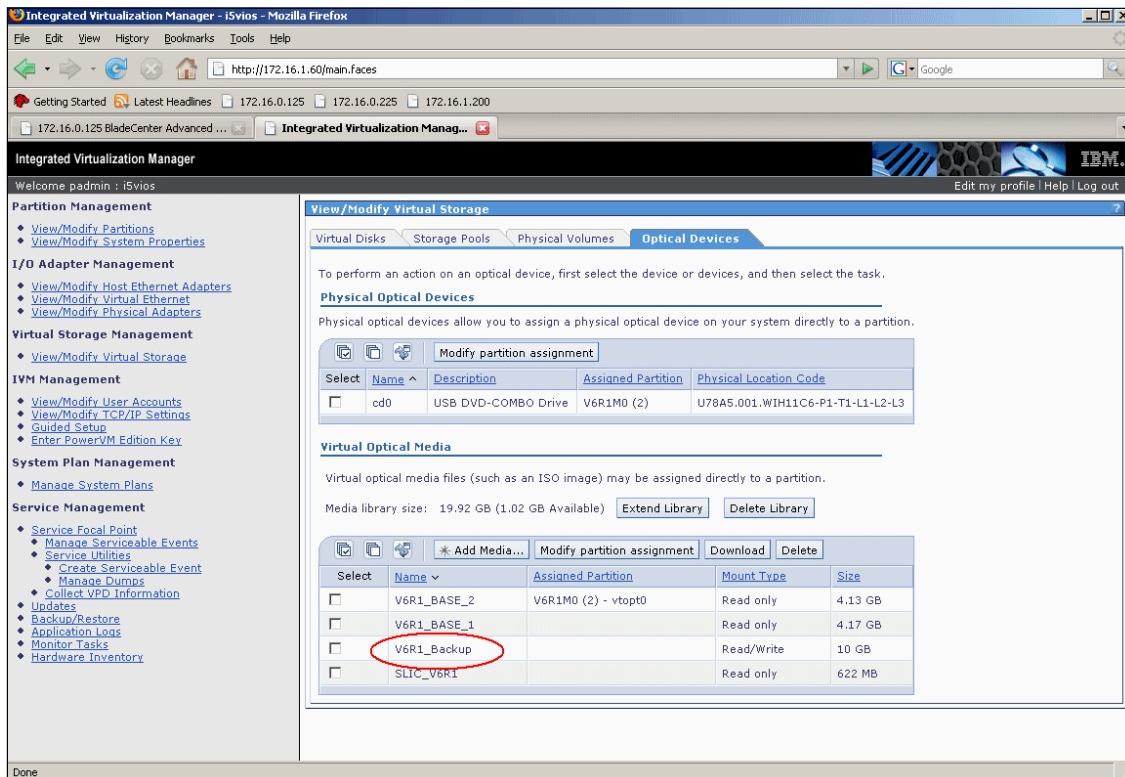


Figure 7-71 IVM Virtual optical device created

To assign the new created virtual optical device to the IBM i 6.1 partition, select the virtual optical device and click **Modify partition assignment** as shown in Figure 7-72 on page 320.

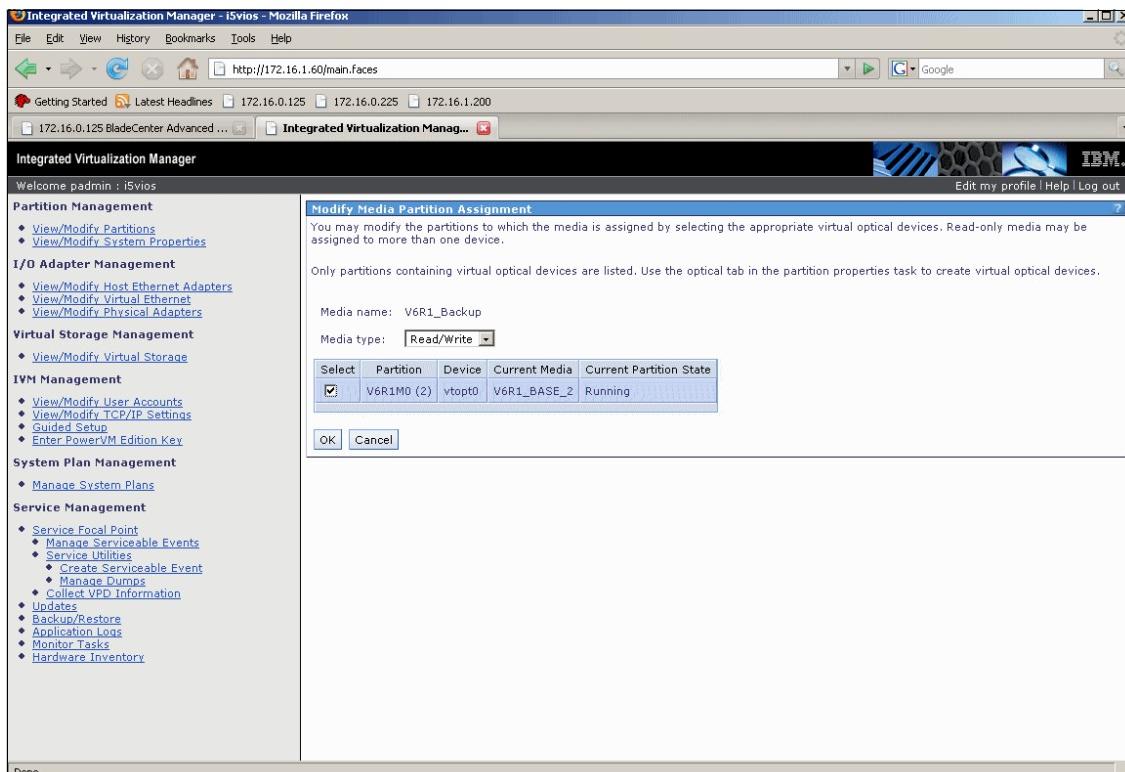


Figure 7-72 Virtual optical device assign to partition

9. Select the IBM i 6.1 partition and click **OK** to continue.

Figure 7-73 on page 321 shows the IVM Virtual Storage Management window with the current assignment of the virtual optical device to the partition.

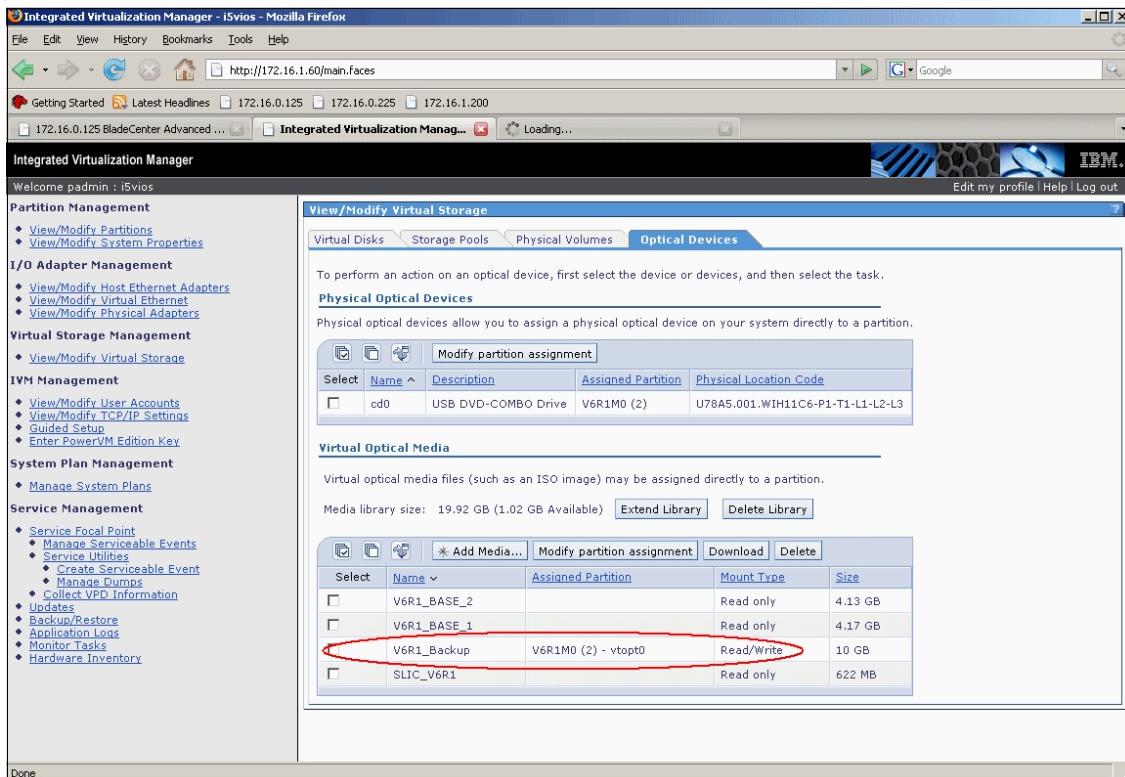


Figure 7-73 IVM Virtual optical device assignment done

After the virtual optical device is mounted to the correct virtual optical device, it will become available in the IBM i 6.1 partition. The IBM i 6.1 Operating System will not use the device name of the virtual optical device given in Integrated Virtualization Manager.

An IBM i 6.1 screen execute command **WRKOPTVOL** and the screen shown in Figure 7-74 on page 322 should appear. The virtual optical device will be identified with a time stamp volume ID.

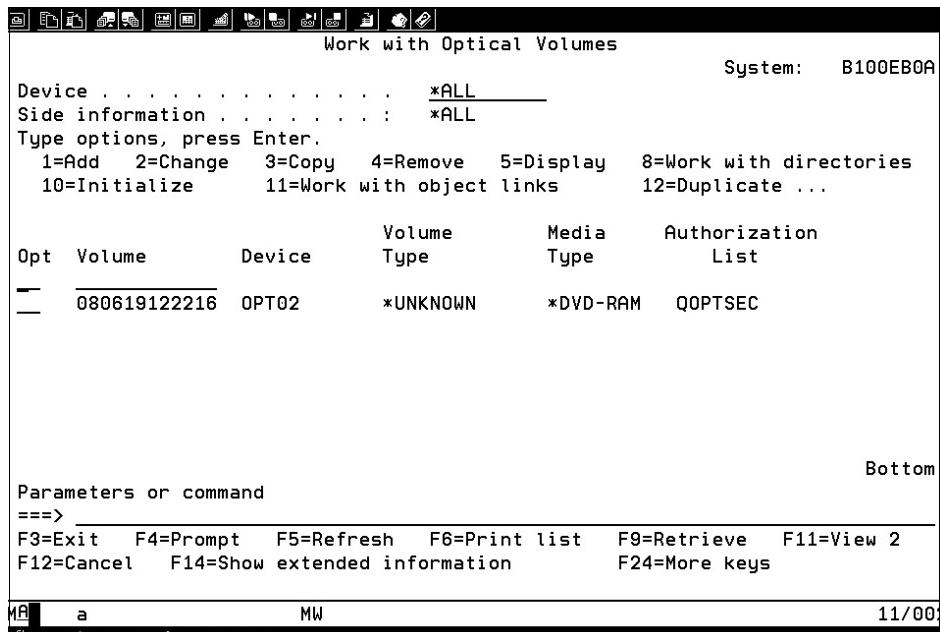


Figure 7-74 Virtual optical device check device

7.5.2 Creating Virtual Media Library using IVM

This section describes the process to create a virtual media library using IVM. This library is created using the IVM options and is located in the /var/vio/VMLibrary directory. Once the library has been created you can add files such as iso images to perform installations of partitions.

1. To begin, a storage pool needs to be created to contain the virtual optical library. If a usable storage pool exists you can skip to step 8 on page 323.
 2. Use the IVM options to select **View/Modify virtual storage**.
 3. Select the **Storage Pools** tab.
 4. Click on **Create Storage Pool**.

Figure 7-75 on page 323 shows an example of the Storage Pools tab highlighting the Create Storage Pool option.

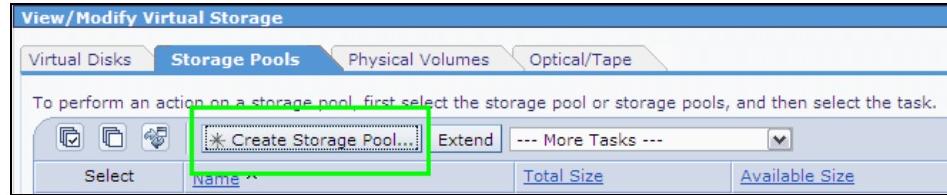


Figure 7-75 Create Storage Pool option

5. Provide a storage pool name.
6. Select the option **Logical Volume Based** for storage pool type.
7. Select one of the available hdisk resources to create the storage pool on.

Figure 7-76 on page 323 provides an example of the storage pool name, size and hdisk selection.

Select	Physical Volume	Size	Physical Location Code
<input type="checkbox"/>	hdisk9	15 GB	U78A5.001.WIH23EC-P1-C11-L1-T1-W500507630E87FE3F-L4010401F000000
<input checked="" type="checkbox"/>	hdisk2	30 GB	U78A5.001.WIH23EC-P1-C11-L1-T1-W500507630E87FE3F-L40104006000000
<input type="checkbox"/>	hdisk3	30 GB	U78A5.001.WIH23EC-P1-C11-L1-T1-W500507630E87FE3F-L40104007000000
<input type="checkbox"/>	hdisk4	30 GB	U78A5.001.WIH23EC-P1-C11-L1-T1-W500507630E87FE3F-L40104008000000
<input type="checkbox"/>	hdisk5	20 GB	U78A5.001.WIH23EC-P1-C11-L1-T1-W500507630E87FE3F-L4010400D000000
<input type="checkbox"/>	hdisk6	20 GB	U78A5.001.WIH23EC-P1-C11-L1-T1-W500507630E87FE3F-L4010400E000000
<input type="checkbox"/>	hdisk7	20 GB	U78A5.001.WIH23EC-P1-C11-L1-T1-W500507630E87FE3F-L4010400F000000
<input type="checkbox"/>	hdisk8	20 GB	U78A5.001.WIH23EC-P1-C11-L1-T1-W500507630E87FE3F-L40104010000000
<input type="checkbox"/>	hdisk10	15 GB	U78A5.001.WIH23EC-P1-C11-L1-T1-W500507630E87FE3F-L40104020000000
<input type="checkbox"/>	hdisk11	15 GB	U78A5.001.WIH23EC-P1-C11-L1-T1-W500507630E87FE3F-L40104021000000
<input type="checkbox"/>	hdisk12	15 GB	U78A5.001.WIH23EC-P1-C11-L1-T1-W500507630E87FE3F-L40104022000000
<input type="checkbox"/>	hdisk13	15 GB	U78A5.001.WIH23EC-P1-C11-L1-T1-W500507630E87FE3F-L40104023000000

* Required field

Figure 7-76 Selecting storage pool name, size and resource

Now that the storage pool has been created the virtual media library can be created using the new storage pool.

8. Use the IVM options to select **View/Modify virtual storage**.
9. Click on the **Optical/Tape** tab.

10. Expand the section **Virtual Optical Media**.

11. Click on **Create Library**.

Figure 7-77 on page 324 shows an example of the Create Library option.

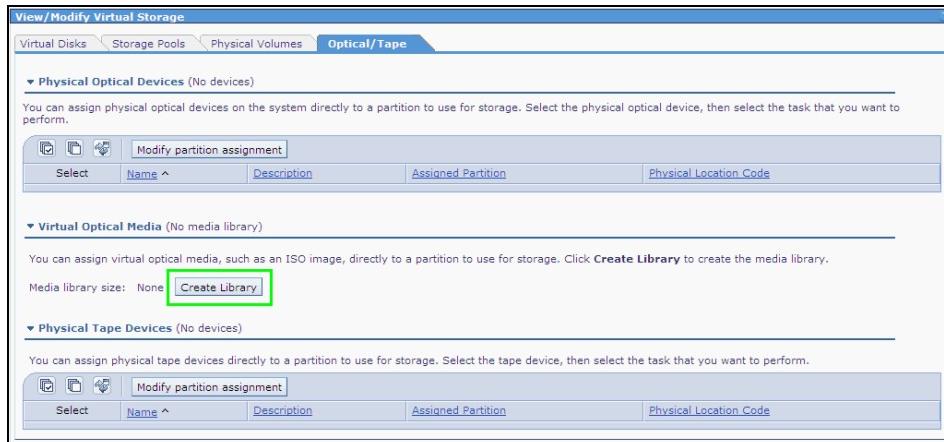


Figure 7-77 *Create Media Library*

12. Define the media library size.

Figure 7-78 on page 324 shows an example of the storage pool name field. Select the correct storage pool to contain your virtual media library.

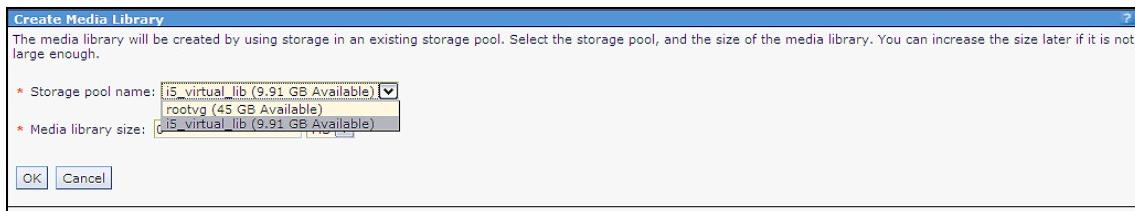


Figure 7-78 *Select storage pool name*

13. Click **OK** to finish.

7.5.3 Adding Image Files to Media Library

One other method of adding media files to the library can be accomplished by creating an iso image of the media and using ftp to copy the file to the JS23/JS43. If your media library has not been created yet, use the section named "Creating a virtual media library for backup" on page 316.

For example: A new installation of IBM i OS in an IBM i partition.

- ▶ Create ISO image files of the installation media.
 - Load the IBM i SLIC media in your PC CDROM
 - Using Record Now or another burning program, create an ISO image of the CD. Usually this is performed using a backup function.

The next few graphics provide an example of using Record Now to create an ISO image of your media.

Figure 7-79 on page 325 shows the option to Save Image. The option is found under the *backup projects* task

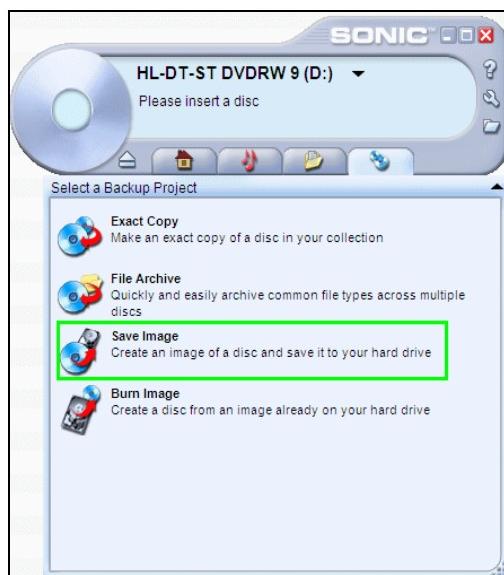


Figure 7-79 Save Image - Create ISO file

Figure 7-80 on page 326 shows an example of selecting the source and destination folders. Click **Browse** to define the output folder.

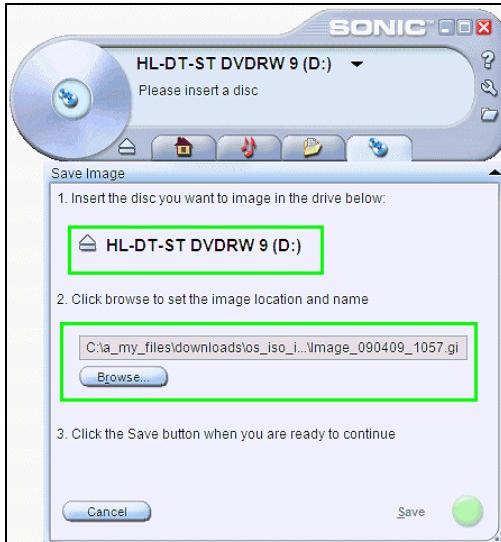


Figure 7-80 Select the output destination folder

Figure 7-81 on page 326 shows an example of the destination folder. Select the **Save as Type** option and ensure the type is set for ISO. It is not the default so it should be changed.

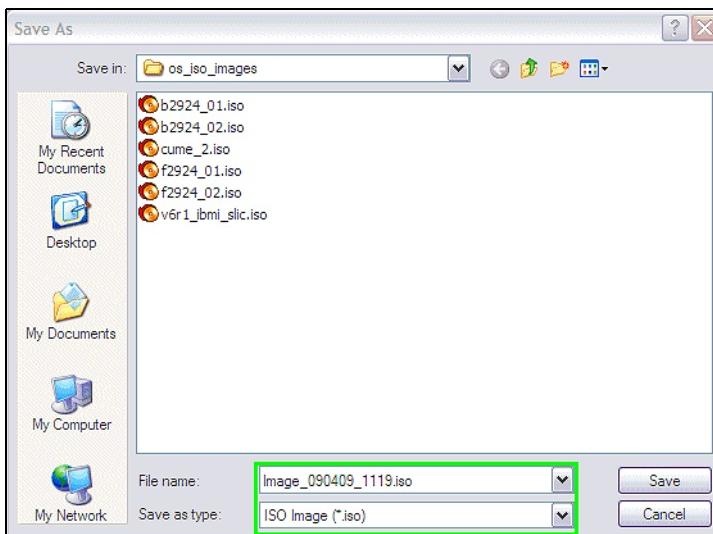


Figure 7-81 Change file type to .iso

- ▶ Copy the iso image file to the JS23/JS43 using ftp. The file will be copied to the /home/padmin directory. Make sure to use image mode when copying the file with ftp. This transfers the file in binary format.
- ▶ Move (*mv*) the .iso file from /home/padmin to /var/vio/VMLibrary. You will need to use *oem_setup_env* to escape the VIOS restricted shell environment to be able to use the *mv* command. It is also recommended to change the file name so it is easier to identify the files. For example: file name ibmi_slic.iso would indicate that the image is the Licensed Internal Code cd needed for a D-IPL.

7.5.4 Attaching a remote PC file or Media device

Using the AMM you can attach a remote PC file or the remote PC media device to the AMM for use as an input device for files. This feature creates a CD device under the AMM for assignment to one of the blade servers. The assignment is performed through the media tray of the bladecenter. After creating the link between the AMM and the remote PC you must assign the blade center media tray to the blade server you are intending to use the file on.

Once the media tray has been assigned to the blade you can use the View/Modify Virtual Storage task and the Optical/Tape tab to assign the CD device to the partition you are working with. The remote PC file and/or CDROM device will appear as an additional CD device assignable to a partition.

The next few graphics will demonstrate how to attach the remote PC file and CDROM. Once the file or CD has been assigned, the device or file can be used by one of the blade partitions by modifying the partition assignment.

1. Open a session to the AMM you wish to work with.
2. Select the **Blade Tasks** option.
3. Select the **Remote Control** option.

Figure 7-82 on page 328 shows an example of the AMM interface and the blade tasks selected.

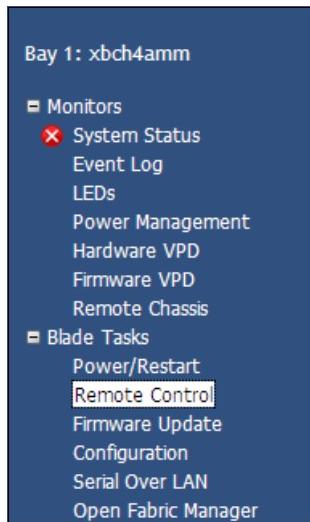


Figure 7-82 Blade Task - Remote Control

4. Once the task Remote Control has been selected, use the **Start Remote Control** button. This will invoke a Java™ window.

Figure 7-83 on page 328 shows an example of the Start Remote Control button.

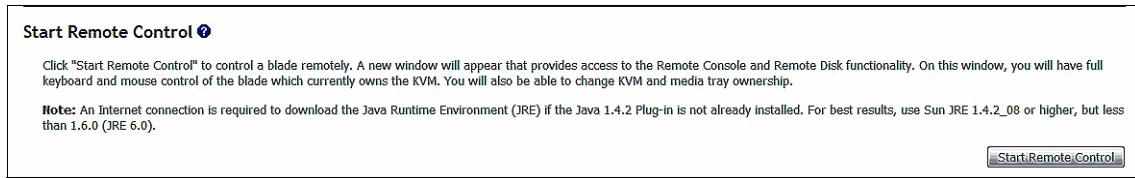


Figure 7-83 Start Remote Control

5. Once the java interface has started, select the **Remote Drive** option.

Figure 7-84 on page 328 shows an example of the java interface for remote control.

Note: It is recommended to leave the Java interface window running until you have completed the media installation action.



Figure 7-84 Remote Control Java window

6. After selecting the Remote Drive option, you will see the Remote Disk window appear. Select the **CD ROM** and/or **Select Image** option. You can use either or both. Figure 7-85 on page 329 provides an example of the Select Image option.

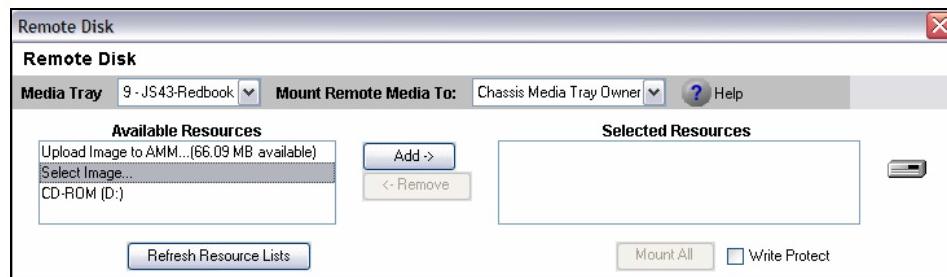


Figure 7-85 Select image option

7. Select the **Add** button. You will then be able to browse for the specific file you want to add as shown in Figure 7-86 on page 329.

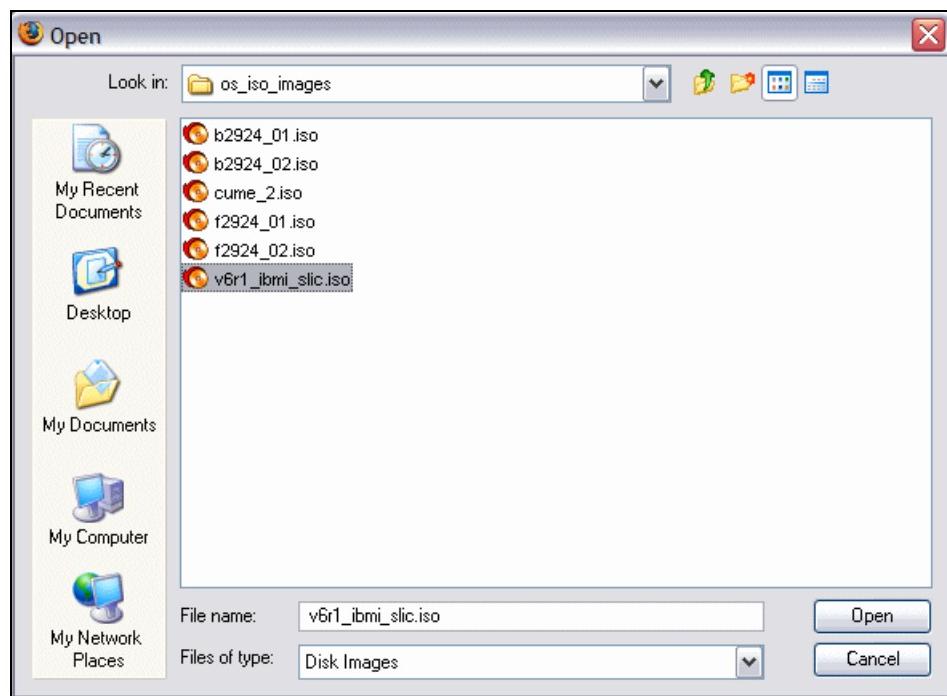


Figure 7-86 Browse and select file

After the file has been added it will appear under the Selected Resources list. Figure 7-87 on page 330 provides an example of this view.

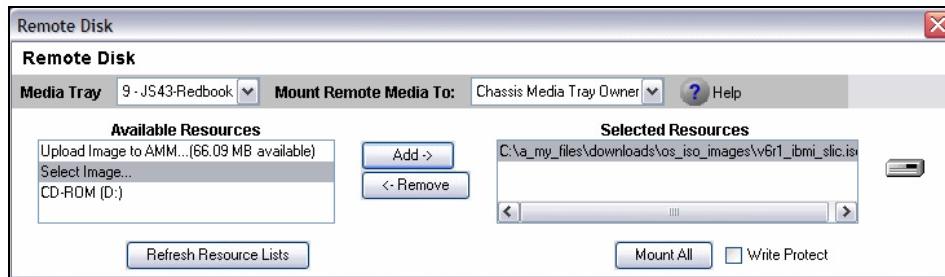


Figure 7-87 File added to Selected Resources list

8. To add the CDROM, select the CDROM listed and click the **Add** button. It will then be listed under the Selected Resources list.
9. After all selections have been made click on the Mount all button. This will add your resources to the AMM and make them available to the blade that has the media tray selected.

Figure 7-88 on page 330 shows an example of the panel after the devices have been mounted. Using the *Unmount All* option would remove the devices from the media tray, however be sure the partition assignment is clear before you remove the resource using the *Unmount* option.

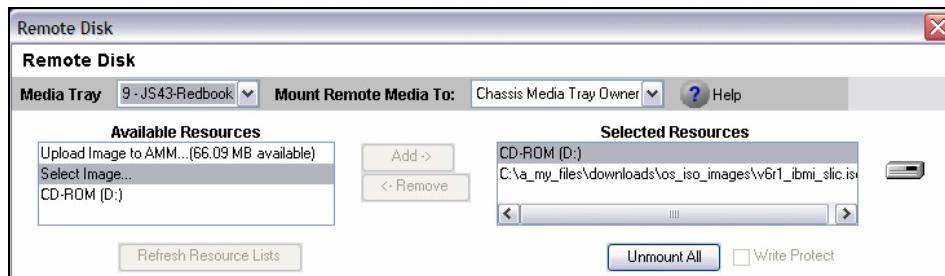


Figure 7-88 Selected Resources mounted

10. After mounting the resources and assigning the media tray to the blade, you will need to run the **cfgdev** command to complete the assignment of the media to the blade. This command is run from the command line of the blade that has the media tray assignment.
11. Once **cfgdev** has completed, use the **View/Modify Virtual storage** task. Then select the **Optical/Tape** tab. The new resources will appear as physical optical devices. Modify the partition assignment for the device to add it into the partition configuration for use.

Figure 7-89 on page 331 shows an example of the resources added using the above process.



Figure 7-89 New physical optical devices

7.5.5 IBM Tivoli Storage Manager

Starting with Integrated Virtualization Manager V1.4, you can install and configure the IBM Tivoli® Storage Manager (TSM) client on the Virtual I/O Server (VIOS). With IBM Tivoli Storage Manager, you can protect your data from failures and other errors by storing backup and disaster recovery data in a hierarchy of offline storage.

IBM Tivoli Storage Manager can help to protect computers running a variety of different operating environments, including the VIO Server, on a variety of different hardware. Configuring the IBM Tivoli Storage Manager client on the Virtual I/O Server enables you to include the Virtual I/O Server in your standard backup.

The TSM client software is included in the VIO Server install images by default. To ensure the IBM Tivoli Storage Manager client is installed on the VIO Server, run the command shown in Example 7-1. This command is performed outside of the VIOS restricted shell environment. The command `1ssw` can be used inside of the shell to see similar results.

Example 7-1 TSM client check

```
1slpp -L | grep TSM
5.4.0.0    C    F    TSM Client - Application
5.4.0.0    C    F    TSM Client - Backup/Archive
5.4.0.0    C    F    TSM Client - Backup/Archive
5.4.0.0    C    F    TSM Client - IMAGE Backup
....
```

Using the IBM Tivoli Storage Manager does not eliminate the IBM i 6.1 save on a virtual optical media device in IVM.

Providing details of configuring and using the IBM Tivoli Storage Manager client and server is beyond the scope of this book. For detailed information about how to configure and manage the VIO Server as a IBM TSM client, refer to:

http://publib.boulder.ibm.com/infocenter/systems/scope/hw/topic/iphb1/i_phb1tivagents.htm?resultof=%22tivoli%22

For more technical information about integrating IBM Tivoli Storage Manager, refer to *PowerVM Virtualization on IBM System p Managing and Monitoring*, SG24-7590.

7.5.6 IBM i 6.1 shutdown and restart

This section describes the shutdown and restart procedure on an IBM i 6.1 partition.

Shut down an IBM i 6.1 partition

It is not recommended to use the View/Modify Partitions - Shutdown option to shutdown an IBM i partition. This option can be used but it is similar to using the white power off button on a true system. If you decide to use this option, do not select Immediate for the shutdown type.

Before you shut down an IBM i 6.1 logical partition, ensure that all jobs are completed and all applications are ended. The proper way to do this is to end all the subsystems.

The correct way to shut down an IBM i 6.1 logical partition is:

- ▶ Open a Console connection using the System i Access for Windows software.
- ▶ If TCP/IP is already configured, you can use a telnet 5250 session however if you are using Telnet, the session will be dropped when the TCP/IP subsystem and running jobs end. You would need to establish the LAN console session to complete the proper shutdown sequence.

To shut down the IBM i 6.1 partition, follow these steps:

1. Open a console connection to the IBM i partition.
2. Sign in with an user ID that has the privilege to shut down the partition.
3. End the running subsystems by issuing the **ENDSBS *ALL** command.
4. Monitor the message queue by using the command **DSPMSG QSYSOPR**
5. Wait for the message “*System ended to restricted condition*”. It may take a few minutes for this message to appear while various jobs and subsystems end. You may need to refresh the display using F5.

6. Execute the command **PWRDWNSYS** in the command line, then use F4 to prompt for options as shown in Figure 7-90 on page 333. Change the *Controlled end delay time* to 300. Press enter when ready to power down the partition.

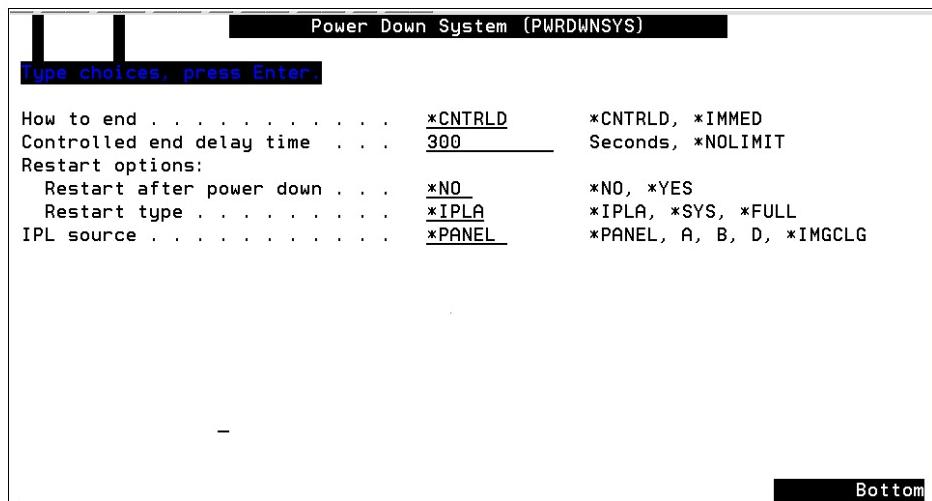


Figure 7-90 IBM i power down partition

7. Confirm the shutdown action by pressing **F16**.
8. This process can take a while. Check the Integrated Virtualization Manager (IVM) window for the message Not Activated in the State column of the IBM i partition.

Start an IBM i 6.1 partition

Before activating a IBM i partition, verify the IPL type you wish to perform is set correctly. Use View/Modify partitions, then select the partition and view the properties. On the General tab you will see the IPL type.

To activate an IBM i 6.1 partition, follow these steps:

1. In Integrated Virtualization Manager (IVM), select the IBM i partition and click **Activate**.
2. Click **OK** in the next window to activate the IBM i partition.
3. After the partition starts its IPL sequence you can activate the Operations Console session to connect to the partition.



Red Hat Enterprise V5.3 Linux installation

This chapter describes the procedures to install Red Hat Enterprise Linux V5.3 on a JS23 BladeCenter.

We discuss the following topics:

- ▶ “Supported Red Hat operating systems” on page 336
- ▶ “Linux LPAR installation using DVD” on page 337
- ▶ “Linux network installation (detailed)” on page 341
- ▶ “Native Red Hat Enterprise Linux 5.3 installation” on page 353
- ▶ “Red Hat Enterprise Linux 5.3 automated installation” on page 354
- ▶ “IBM service and productivity tools” on page 365

8.1 Supported Red Hat operating systems

Red Hat Enterprise Linux for POWER Version 4.6 or later and Red Hat Enterprise Linux for POWER Version 5.1 or later support installation on a JS23.

This chapter specifically covers installing Red Hat Enterprise Linux for POWER Version 5.3 with a DVD and over the network on a PowerVM logical partition (LPAR).

8.1.1 Considerations and prerequisites

There are some system configuration considerations and prerequisites prior to installing Red Hat Enterprise Linux 5.3 on a JS23 partition. They are covered here.

PowerVM LPAR considerations

By the characteristics of the virtualization features in the System p JS23 BladeCenter, the operating system and applications do not know they are running in a micro-partitioned or virtualized I/O environment. This allows applications to run unmodified in a partition that takes advantage of virtualization features.

Because the virtual partition handles the transition of the virtual adapter's I/O operation to the physical adapter, it is important to guarantee that the partition is properly sized to handle the I/O requirements. A good source for processor and memory requirements for PowerVM partitions based on I/O requirements is found at:

<http://www14.software.ibm.com/webapp/set2/sas/f/vios/documentation/perf.html>

Additional information on virtualization and LPAR setup is available in Chapter 4, “System planning and configuration using VIOS with IVM” on page 71.

Red Hat Enterprise Linux 5.3 installation considerations

The minimum RAM required to install Red Hat Enterprise Linux 5.3 is 1 GB, while the recommended RAM is 2 GB.

Note: Make sure you have 1G of RAM or more allocated to your LPAR. Systems with less than 1 GB of RAM may experience installation process hangs or other unexpected failures.

In addition, ensure there is enough unpartitioned disk space or have one or more partitions that can be deleted to free up disk space for the Linux installation. The Red Hat Recommended Partitioning Scheme is available at:

http://www.redhat.com/docs/en-US/Red_Hat_Enterprise_Linux/5.3/html/Installation_Guide/ch11s03.html

8.2 Linux LPAR installation using DVD

With PowerVM installed and the system partitioned into LPARs using the PowerVM LPAR considerations and Red Hat Enterprise Linux 5.3 prerequisites, we are ready to install Linux on the JS23 BladeCenter LPAR.

Start an LPAR installation of Red Hat Enterprise Linux 5.3 using a DVD

The steps to start the installation of Red Hat Enterprise Linux 5.3 from a DVD are as follows:

1. On the Advanced Management Module (AMM) select **Blade Tasks** → **Remote Control**.
2. In the right side window, look for Media Tray Owner and click the drop-down arrow to display the various blade bays. Select your blade bay location with the mouse as shown in Figure 8-1.

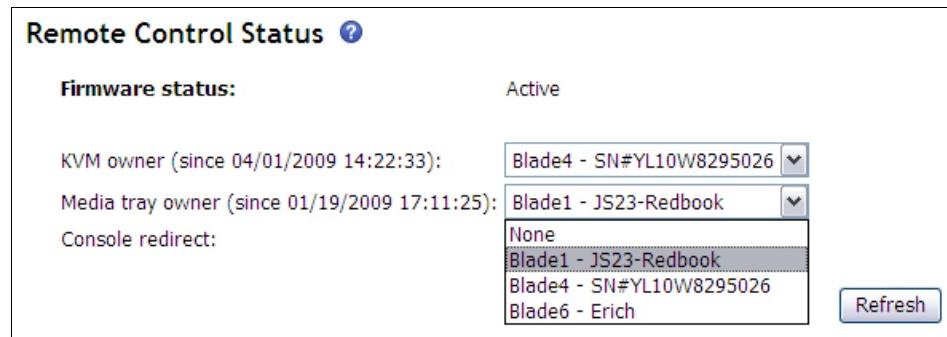


Figure 8-1 Remote Control window - assign Media Tray

3. Press **Refresh** on the right side.

Important: The other option is to press the MT button on the blade to assign the media tray to the blade.

Make sure no other blade in the BladeCenter is using the media tray before pressing this button. Their MT light is on if the media tray is assigned to them.

- Double-check that your blade bay owns the media tray by opening the AMM window and selecting **Monitors** → **System Status**. The right window will show a “check mark” in the MT column of your blade bay location. Figure 8-2 gives one example of this assignment.

Blades ?								
Click the icon in the Status column to view detailed information about each blade.								
Bay	Status	Name	Pwr	Owner **		cKVM*	I/O	Compatibility
				KVM	MT*			
1		JS23-Redbook	On	✓	✓		OK	
2		<i>No blade present</i>						

Figure 8-2 BladeCenter System status.

- Place the DVD into the BladeCenter media tray.
- Log in to the Integrated Virtualization Manager (IVM) via a Web browser.
 - Select one of the available PowerVM LPARs to install with Red Hat Enterprise Linux 5.3 by placing a checkmark in the Select box.
 - Click **Activate**, as shown in Figure 8-3.

Partition Details										
Select	ID ^	Name	State	Uptime	Memory	Processors	Entitled Processing Units	Utilized Processing Units	Reference Code	
<input type="checkbox"/>	1	is23-vios	Running	20.33 Hours	1 GB	4	0.4	0.57		
<input checked="" type="checkbox"/>	2	is23-lp1	Not Activated		1 GB	1	0.1		00000000	

Figure 8-3 Activating an IVM partition

- c. Click the drop-down arrow to the right of the More Tasks field and select **Open terminal window**.

Important: Make sure the latest Java Runtime Environment (JRE™) is installed on the native system to run the IVM terminal. At the time of this publication, the recommended JRE is Sun's JRE 1.4.2_19, or higher.

Note: Even though this section covers installation via the Integrated Virtualization Manager (IVM) console, there are other console options available on the JS23. They are covered in Appendix A, “Consoles, SMS, and Open Firmware” on page 493.

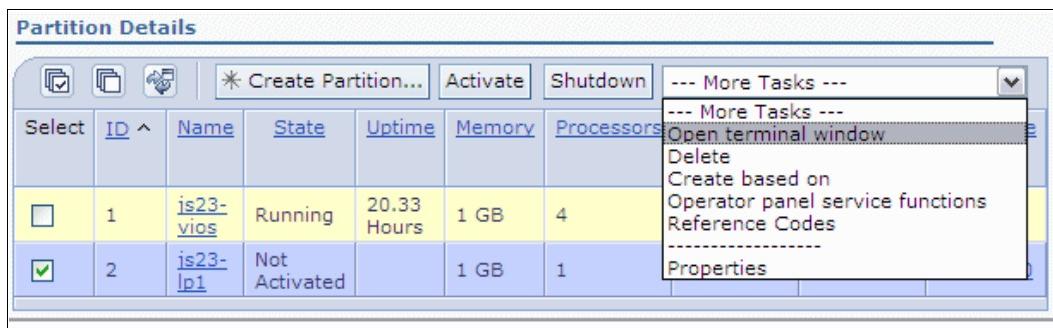


Figure 8-4 Opening a terminal window from the IVM

The console is a pop-up and it will ask you to authenticate with the PowerVM User ID and password.

7. The SMS menu appears in the IVM terminal; see Figure 8-5 on page 340.

Tip: The SMS menu appears very quickly after activating the box, so have your hand ready on the 1 key so you can press it immediately when you see the SMS menu. If you happen to miss the SMS option section window, just shut down and reactivate the LPAR again. The IVM terminal will hold the connection to the PowerVM LPAR even if the LPAR is shut down.

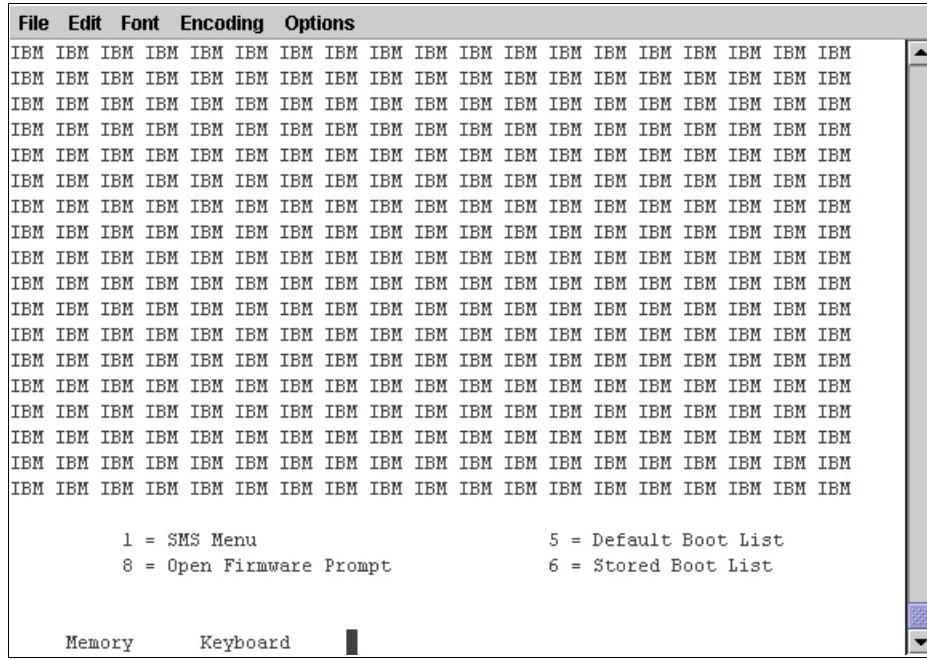


Figure 8-5 SMS menu

- a. Select 1 = SMS Menu by pressing the number 1 on the keyboard.

Tip: Press the number next to the desired system function to select and navigate through the SMS menu.

- b. Select option 5. **Select Boot Options**.
- c. Choose option 1. **Select Install/Boot Device**.
- d. Pick 3. **CD/DVD**.
- e. Select 6. **USB**.
- f. Finally, select 1. **USB CD-ROM**.
- g. Choose 2. **Normal Mode Boot**.
- h. Pick 1. **Yes** to exit the SMS menu.
- i. At the boot: prompt press the Enter key.

At this point the system will begin reading off the DVD, which can take a few minutes.

8. After successfully booting from the DVD, the Red Hat Anaconda Installer will ask if you want to perform a media check. If this is your first time installing

using this media, we highly recommend running the media check. Once the media check is complete, Anaconda will assist with the completion of the install. More detailed installation instructions are available here:

http://www.redhat.com/docs/en-US/Red_Hat_Enterprise_Linux/5.3/html/Installation_Guide/pt-install-info-ppc.html

8.3 Linux network installation (detailed)

This section describes a Network File System (NFS) installation on an external Storage Area Network (SAN) device which is part of a PowerVM LPAR. Even though this example uses a specific storage device and network option to complete the installation, this process is generic enough that it will work for all supported internal or external storage devices on JS23 and only a few slight modifications are needed in Anaconda's setup if a different network installation option is selected. The goal of this section is to show the external SAN disk install in detail so the steps can be referenced later.

This installation is using an NFS installation method, but NFS is one of many supported network installation types on Linux. The NFS server configuration and setup are described in detail in Appendix C, “Additional Linux installation configuration options” on page 535.

This section assumes that you have already set up the NFS server properly and have read “PowerVM LPAR considerations” on page 336 and followed the installation prerequisites in “Red Hat Enterprise Linux 5.3 installation considerations” on page 336.

Start an LPAR installation of Red Hat Enterprise Linux 5.3 over the network

To start an LPAR installation of Red Hat Enterprise Linux 5.3 over the network, use the following steps:

1. Log in to the Integrated Virtualization Manager (IVM) using a Web browser.
 - a. Select one of the available PowerVM LPARs to install with Red Hat Enterprise Linux 5.3 by placing a checkmark in the Select box.

- b. Click **Activate**, as shown in Figure 8-6.

Select	ID ^	Name	State	Uptime	Memory	Processors	Entitled Processing Units	Utilized Processing Units	Reference Code
<input type="checkbox"/>	1	is23-vios	Running	20.33 Hours	1 GB	4	0.4	0.57	
<input checked="" type="checkbox"/>	2	is23-lp1	Not Activated		1 GB	1	0.1		00000000

Figure 8-6 Activating an IVM partition

- c. Click the drop-down arrow to the right of the More Tasks field (Figure 8-7) and select **Open terminal window**.

Select	ID ^	Name	State	Uptime	Memory	Processors	Entitled Processing Units	Utilized Processing Units	Reference Code
<input type="checkbox"/>	1	is23-vios	Running	20.33 Hours	1 GB	4	0.4	0.57	
<input checked="" type="checkbox"/>	2	is23-lp1	Not Activated		1 GB	1	0.1		00000000

Figure 8-7 Opening a terminal window from the IVM

The console is a pop-up and it will ask you to authenticate with the PowerVM User ID and password.

2. The SMS menu appears in the IVM terminal (Figure 8-8 on page 343).

Tip: The SMS menu appears very quickly after activating the LPAR, so have your hand ready on the 1 key so you can press it immediately when you see the SMS menu. If you happen to miss the SMS selection window, just shut down and reactivate the LPAR again. The IVM terminal will hold the connection to the PowerVM LPAR even if the LPAR is shut down.

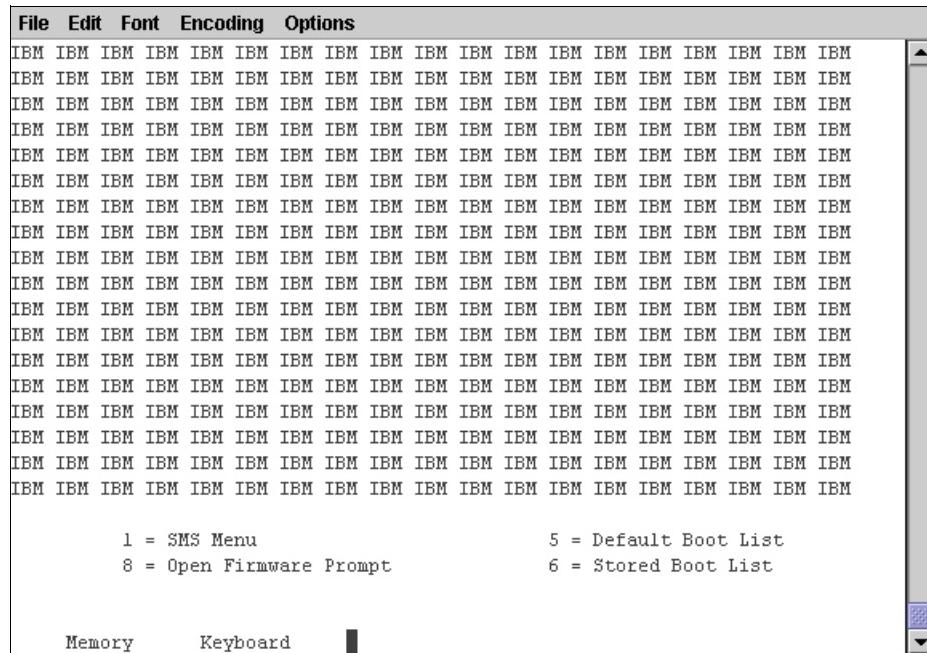


Figure 8-8 SMS menu

- a. Select **1 = SMS Menu** by pressing the 1 key on the keyboard.

Tip: Press the number next to the desired system function to select and navigate through the SMS menu.

- b. Choose **5. Select Boot Options**.
- c. Choose **1. Select Install/Boot Device**.
- d. Choose **6. Network**.
- e. Pick a specific network port.
- f. Choose **2. Normal Mode Boot**
- g. Pick **1.Yes** to exit the SMS menu.

If everything is set up correctly, the system will receive an IP address from the TFTP server and start loading the boot image for the first stage of the install.

Note: Review Appendix C, “Additional Linux installation configuration options” on page 535 prior to starting an NFS installation.

You will notice the packet count value increasing. After the complete boot image is uploaded, the system boots off of it to show the Red Hat Enterprise Linux 5.3 welcome screen, shown in Figure 8-9.

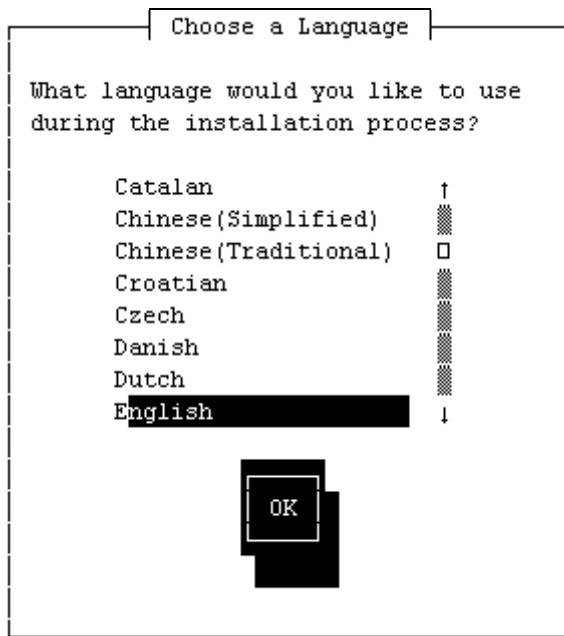


Figure 8-9 Red Hat Enterprise Linux 5.3 welcome screen

3. Select the language to use during the install process. In this example we are using English. Then press the Tab key to move to the **OK** button and then press the Space bar to confirm.

Tip: Use the Tab key to navigate between selections, the Up/Down Arrow keys to move within a list, and the Space bar to select a specific option inside the Anaconda installer.

4. The Installation method is NFS. Use the arrow key to scroll down to NFS image and then press the Tab key to navigate to the **OK** button and press the Space bar to confirm.
5. Select the networking device to install the image from, as shown in Figure 8-10 on page 345

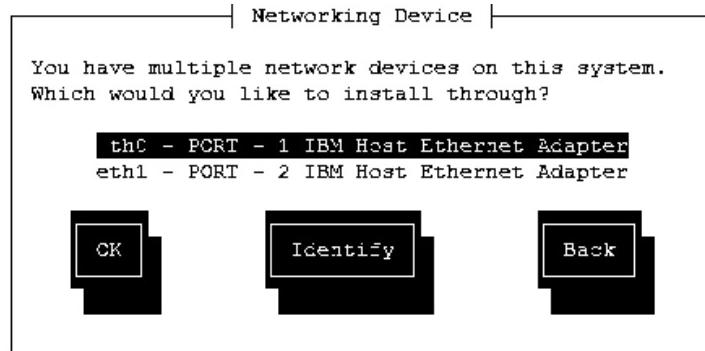


Figure 8-10 Select network device

Note: This step appears only when running Anaconda on machines with more than one network card. The **Identify** option can be used to find the physical port for the selected interface, by flashing the LED lights of the correspondent physical port for a number of seconds.

6. To configure DHCP, select either IPv4 or IPv6 support and then Dynamic IP configuration (DHCP) from the TCP/IP window. Then select **OK**. See Figure 8-11 on page 346 for more details and skip steps 7 and 8.

As an example, if you wish to configure static IP parameters with IPv4 support, configure the TCP/IP parameters and **Enable IPv4 support**, then select **Manual configuration** and disable **Enable IPv6 support** as shown in Figure 8-11, and configure the device as shown in Figure 8-12 on page 346.

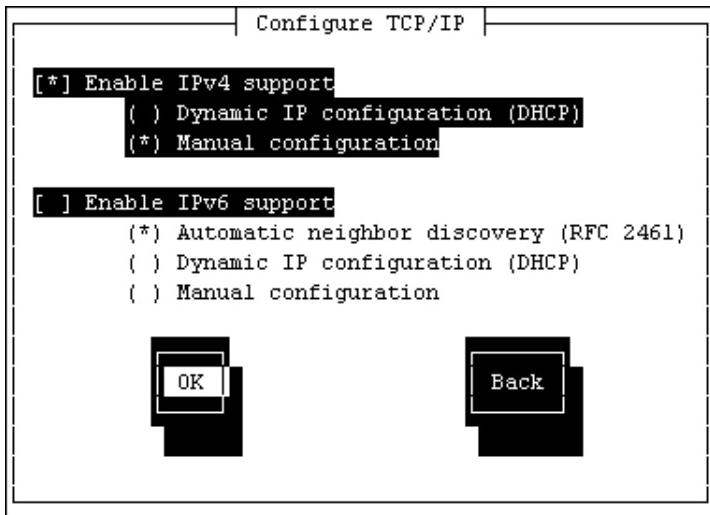


Figure 8-11 TCP/IP configuration panel

7. In the next panel, configure the LPAR's IPv4 address, subnet mask, gateway, and name server. An example configuration is shown in Figure 8-12

A screenshot of a terminal window titled "Manual TCP/IP Configuration". It contains instructions: "Enter the IPv4 and/or the IPv6 address and prefix (address / prefix). For IPv4, the dotted-quad netmask or the CIDR-style prefix are acceptable. The gateway and name server fields must be valid IPv4 or IPv6 addresses." Below the instructions are three input fields: "IPv4 address: xxx.yyy.mm / zzz.zzz.zyx", "Gateway: xx.yy.xxx", and "Name Server: xx.yy.zz". At the bottom are two buttons: "OK" on the left and "Back" on the right.

Figure 8-12 TCP/IP configuration of IP address, gateway, and name server

8. In the NFS Setup window in Figure 8-13 on page 347, enter the IP address of the NFS server and in the field directly below that, enter the NFS directory that contains the Red Hat Enterprise Linux 5.3 install image.
- 9.

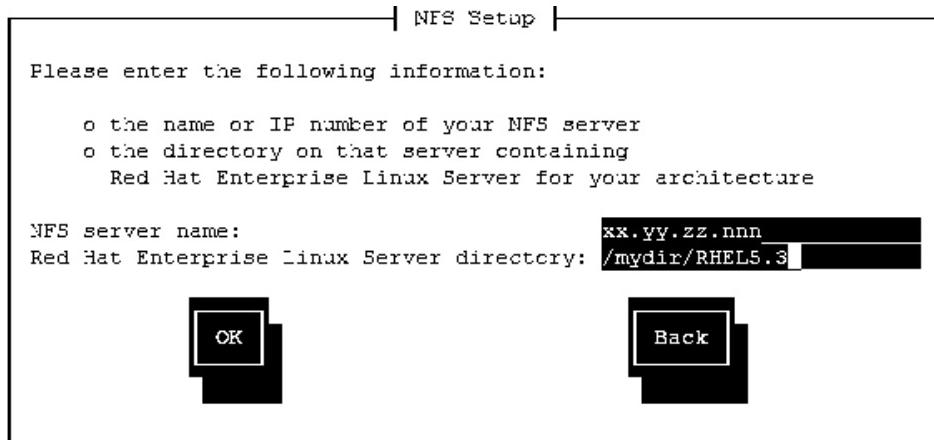


Figure 8-13 NFS server configuration window panel

10. In this step it is possible to start a *Virtual Network Computing* (VNC) server and continue the installation from Anaconda's graphical interface, but for this example we'll continue with the text mode interface, as shown in Figure 8-14.

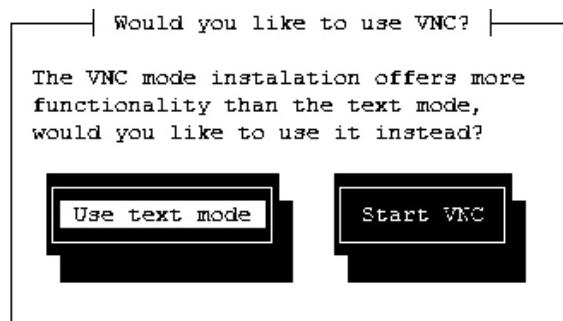


Figure 8-14 Select between VNC or text installation modes panel

11. Approximately one minute later the Welcome to Red Hat Enterprise Linux Server message panel appears. Select **OK**.
12. Enter the installation number in Figure 8-15 on page 348. Click **OK** to confirm.

Note: The installation number is specific to systems running Red Hat Enterprise Linux 5 or later. This number comes from Red Hat subscription management or the poster included in the media kit, and it determines the packages available to the installer.

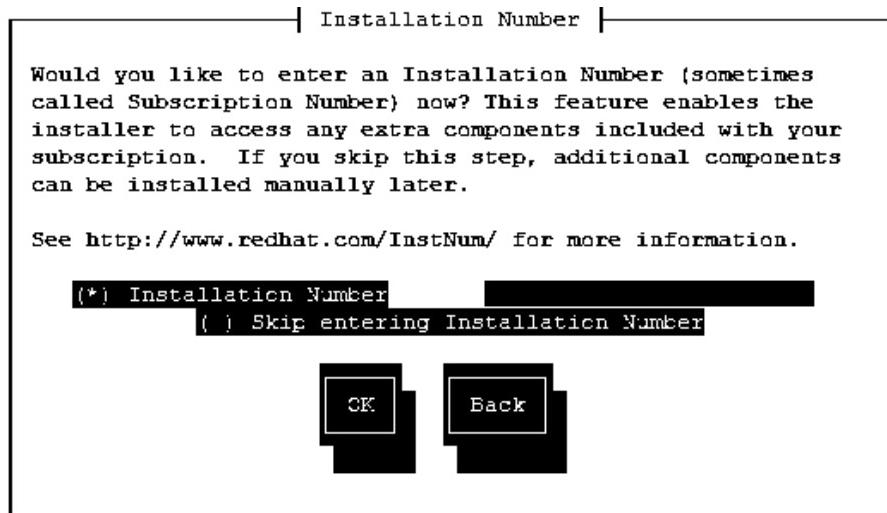


Figure 8-15 Installation number panel

Note: If you skip entering the Installation number, then you will only have the basic packages to select from later on. In this case, a warning will be presented and you'll need to select **Skip** to proceed.

13. Select the disk partitioning type for this installation. In this scenario, we have selected the option **Remove all partitions on selected drives and create a default layout**. The drive we have selected is /dev/sda, which is the externally connected storage SAN disk that was partitioned using the IVM tools. See Figure 8-16 on page 349 as an example.

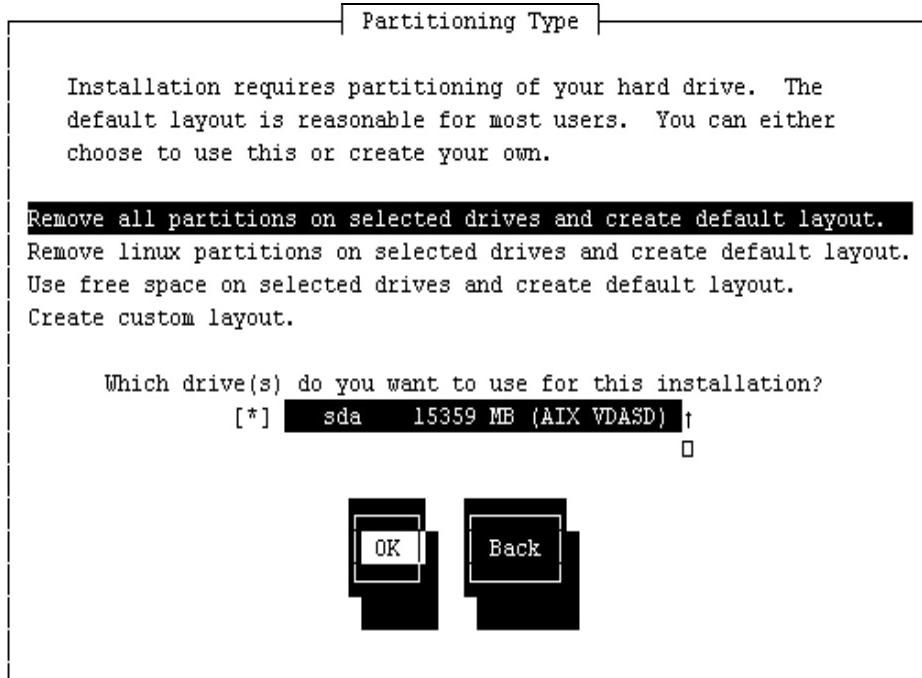


Figure 8-16 Select Partitioning Type panel

14. A warning appears asking if the selection is OK. Press **Yes** to confirm.
15. Select **Yes** to review the suggested disk partition layout.
16. Review the allocated size for swap, ext3 file system, and /boot, as shown in Figure 8-17 on page 350. Press **OK** to confirm.

Partitioning					
Device	Start	End	Size	Type	Mount Point
VG VolGroup00			15232M	VolGroup	
LV LogVol01			1984M	swap	
LV LogVol00			13248M	ext3	/
/dev/sda					
sda1	1	1	0M	PPC PReP B	
Free space	1	1	7M	Free space	
sda2	2	13	94M	ext3	/boot
sda3	14	1958	15257M	physical v	

Figure 8-17 Review Partitioning panel

Note: This configuration can only be edited by a graphical installer such as Virtual Network Connection (VNC). This cannot be done from the IVM terminal, so only the default values selected by the Anaconda Installer are allowed.

17. Press **OK** on the Network Configuration panel. The default is fine because this was already set up in Figure 8-12 on page 346.
18. Press **OK** for the Miscellaneous Network Setting window. The gateway and primary DNS are already configured.
19. Press **OK** after editing/confirming the hostname.
20. Select the time zone for the LPAR using the Tab key to enter the time zone options field and use the arrow keys to move up or down inside of it until you find your locale. Press the Tab key and then Space bar on the **OK** button to confirm.
21. Enter the root password.
22. Select any additional software applications with the Space bar and press **OK** when complete, as shown in Figure 8-18 on page 351

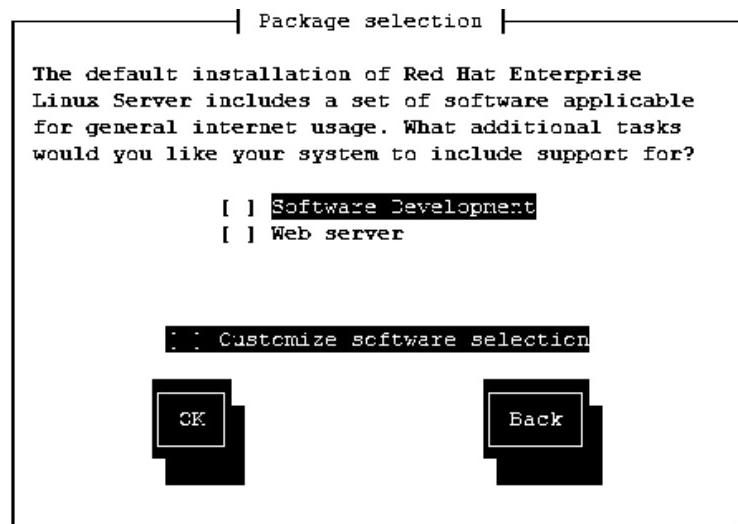


Figure 8-18 Select additional packages panel

Note: These packages can be installed later using `yum` from the command line if you skip this step during the installation.

23. Press **OK** to allow the installation to begin.

The next window has two progress bars: One for the package currently being installed and another detailing the overall progress of the installation.

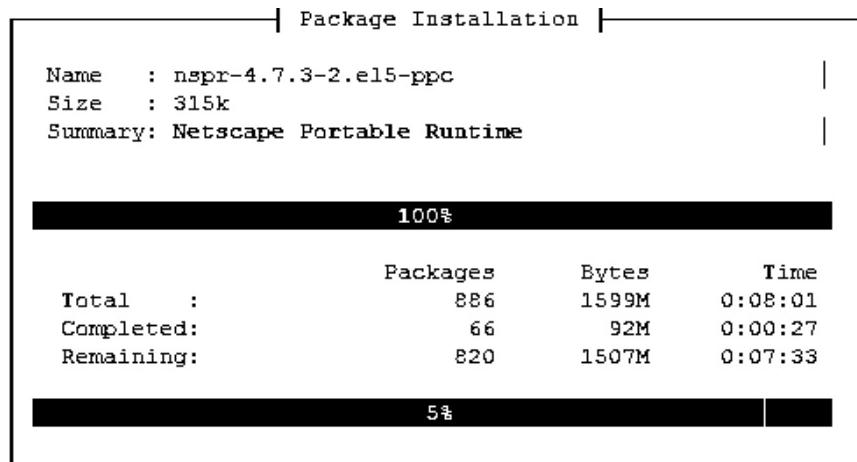


Figure 8-19 Installation progress window

24. Press **Reboot** after the Install Complete window appears, as shown in Figure 8-20

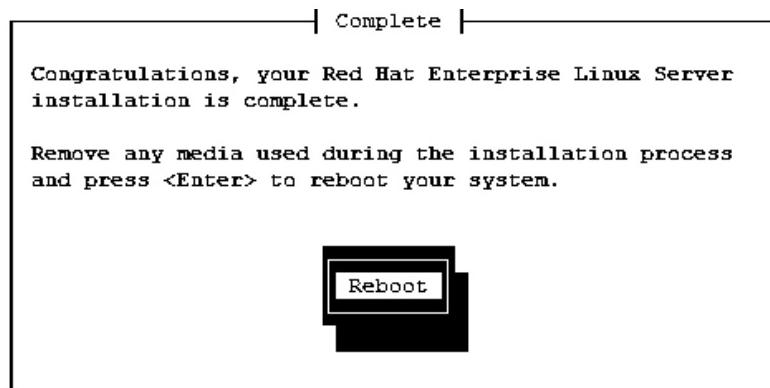


Figure 8-20 Installation complete panel

Note: If the LPAR does not automatically boot from the intended hard disk (boot device) after reboot, try this:

- a. Shut down and reactivate the LPAR from the IVM.
- b. Enter the SMS Menu.
- c. Select **5. Select Boot Options** → **1. Select Install/Boot Device** → **5. Hard Drive** → **9. List All Devices**.
- d. Choose the appropriate hard disk with the Linux image from the given list.
- e. Select **2. Normal Mode Boot** → **1. Yes**.

25. During boot the Setup Agent window appears (Figure 8-21). You can modify any of the fields if desired or press **Exit** to finish booting the LPAR.

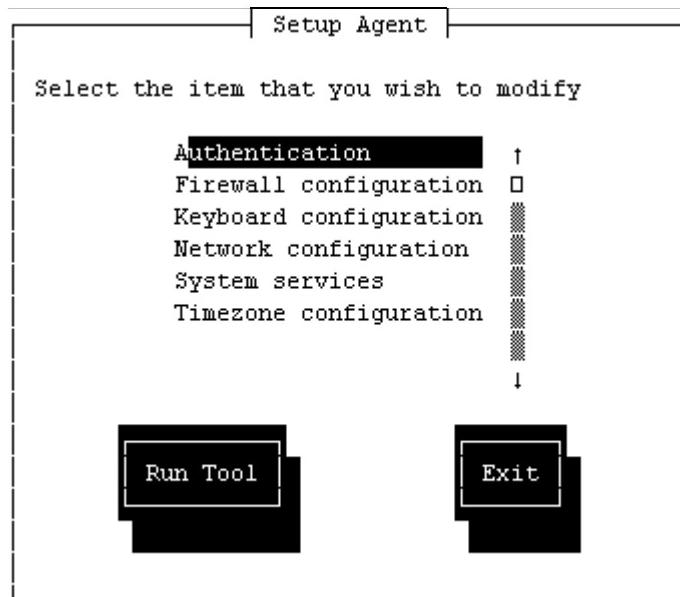


Figure 8-21 Setup Agent panel

The Red Hat Enterprise Linux 5.3 login prompt appears, as shown in Figure 8-22. The installation is complete.

```
[ OK ]
Starting smartd: [ OK ]

Red Hat Enterprise Linux Server release 5.3 (Tikanga)
Kernel 2.6.18-128.el5 on an ppc64

localhost.localdomain login: █
```

Figure 8-22 Finished Red Hat Enterprise Linux 5.3 installation

8.4 Native Red Hat Enterprise Linux 5.3 installation

A native Red Hat Enterprise Linux 5.3 installation on a IBM BladeCenter JS23 Express follows a similar process to the one given in the previous VIOS LPAR installation sections. However, there are some key differences:

- ▶ In a native installation, the IVM terminal is no longer available to complete the Linux installation, but you can use the Serial Over LAN (SOL) console, or the

graphical display (via Blade Center's KVM), as an alternative. See Appendix A, “Consoles, SMS, and Open Firmware” on page 493 for more information. Use the SOL console to display the SMS menu and the Anaconda options during the installation.

- ▶ The resource allocation of processors, I/O adapters, memory, and storage devices in a native environment is *fixed*.
- ▶ Virtualization functions and features are not available.

8.5 Red Hat Enterprise Linux 5.3 automated installation

Red Hat provides an automated installation functionality known as Kickstart. The system administrator performs a Kickstart automated installation by creating a single file containing answers to all the questions normally asked during a Red Hat installation. This file can reside on a single server system, and multiple clients can read it during installation.

This section explains how to create Kickstart files and use them to perform an unattended installation of Red Hat Enterprise Linux 5.3

8.5.1 Kickstart profile creation methods

There are three methods to create a Kickstart file:

1. Create the file using the Red Hat Kickstart Configurator tool.
2. Have Red Hat Anaconda Installer generate this file for you during a manual installation. After the installation completes, the Kickstart file resides at /root/anaconda-ks.conf and you can copy it to an installation server. The basic file created during a manual NFS installation is given in Example 8-1.
3. Generate the file yourself using a text editor.

Example 8-1 Basic Kickstart configuration file created during manual installation

```
# Kickstart file automatically generated by anaconda.

install
nfs --server=9.3.80.16 --dir=/install/linuxRHEL5.3
key 2515-dd4e-2152-25dd
lang en_US.UTF-8
network --device eth1 --bootproto static --ip 9.3.233.177 --netmask
255.255.254.0 --gateway 9.3.232.1 --nameserver 9.3.192.21 --hostname
JS23-turkey-3-1p1.austin.ibm.com
rootpw --iscrypted $1$oVjuAHA9$6cqrT/Nfnh6pRYFBIBECy/
```

```
firewall --enabled --port=22:tcp
authconfig --enableshadow --enablemd5
selinux --enforcing
timezone --utc America/New_York
bootloader --location=partition --driveorder=sda --append="console=hvc0
rhgb quiet"
# The following is the partition information you requested
# Note that any partitions you deleted are not expressed
# here so unless you clear all partitions first, this is
# not guaranteed to work
#clearpart --all --drives=sda
#part prepboot --fstype "PPC PReP Boot" --size=4 --ondisk=sda
#part /boot --fstype ext3 --size=100 --ondisk=sda
#part pv.14 --size=0 --grow --ondisk=sda
#volgroup VolGroup00 --pesize=32768 pv.14
#logvol / --fstype ext3 --name=LogVol00 --vgname=VolGroup00 --size=1024
--grow
#logvol swap --fstype swap --name=LogVol01 --vgname=VolGroup00
--size=1000 --grow --maxsize=1984

%packages
@office
@editors
@text-internet
@gnome-desktop
@dialup
@core
@base
@games
@java
@legacy-software-support
@base-x
@graphics
@printing
@sound-and-video
@admin-tools
@graphical-internet
emacs
kexec-tools
device-mapper-multipath
hfsutils
xorg-x11-utils
xorg-x11-server-Xnest
libsane-hpaio
```

Notice that all of the partition information is commented out with a # symbol. This section needs to be uncommented and edited to support the partition schemes of systems that will use the automated Kickstart install process. The automated Kickstart process will not work without these edits.

8.5.2 Create Kickstart file using Kickstart Configurator

In this section, we use the Kickstart Configurator tool with a graphical interface to demonstrate how to create a basic Kickstart text file.

There are a lot of optional settings in this file, but some are mandatory settings or dependencies. It is impossible to cover every configuration option, but we provide a general overview to get you familiar enough with the Kickstart Configurator tool to navigate on your own.

Restriction: It is necessary to have a running X Server to use the administration tools provided by Red Hat, because in most cases the tools use a graphical user interface (GUI).

You can install a *Virtual Network Computer* (VNC) server in your system and export the graphical interface to some other host in the network. RHEL 5.3 distributes the *vnc-server* rpm package, which you can install to get access to VNC client/server functionalities.

Red Hat provides a utility called *system-config-kickstart* to assist with the creation of the configuration file for an unattended installation. However, after you create the Kickstart file with the tool, you may have to make some manual changes to it.

Note: Kickstart is not included in the default software installation. Install the
pykickstart
system-config-kickstart
rpm packages from the installation media to get Kickstart working.

1. After you install the required Kickstart packages, issue the *system-config-kickstart* command from the command line to launch this utility.
2. A window opens showing the Basic Configuration panel. The most important configuration setting is the Target Architecture setting highlighted with a

rectangle in Figure 0-1. It is also important to define a root password to enable SSH login after installation. This password is encrypted in the configuration file.

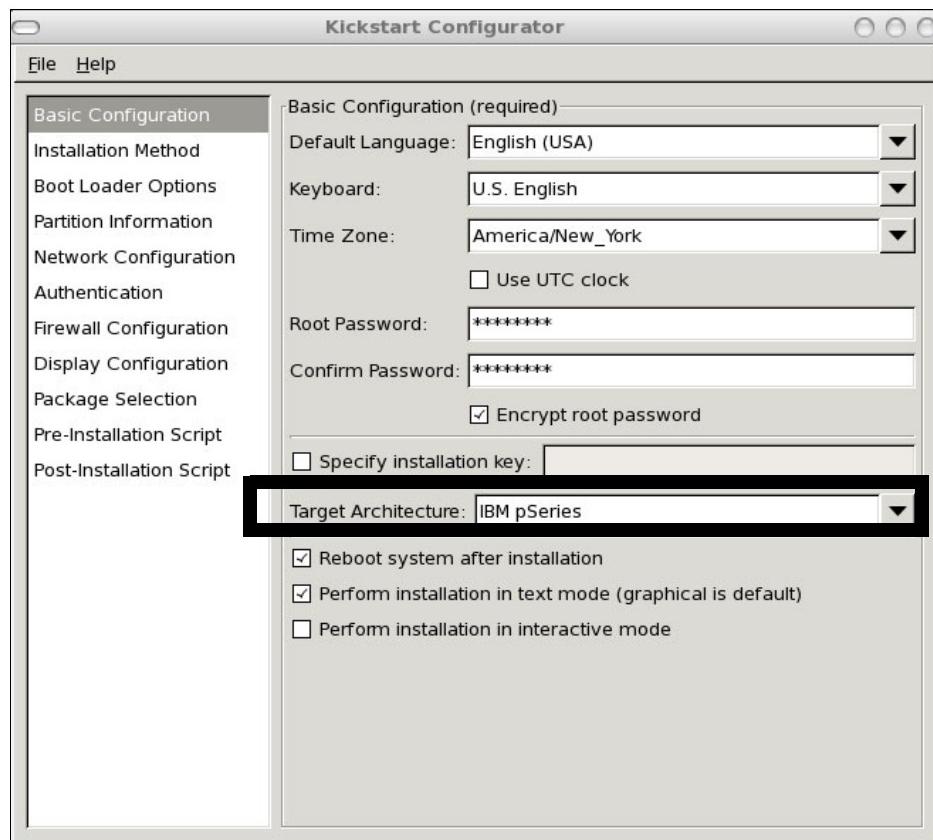


Figure 0-1 Kickstart main window with Basic Configuration panel (©2008 Red Hat, Inc.)

3. In the Installation Method panel (shown in Figure 0-2), all the basic parameters for a network installation using NFS are shown.

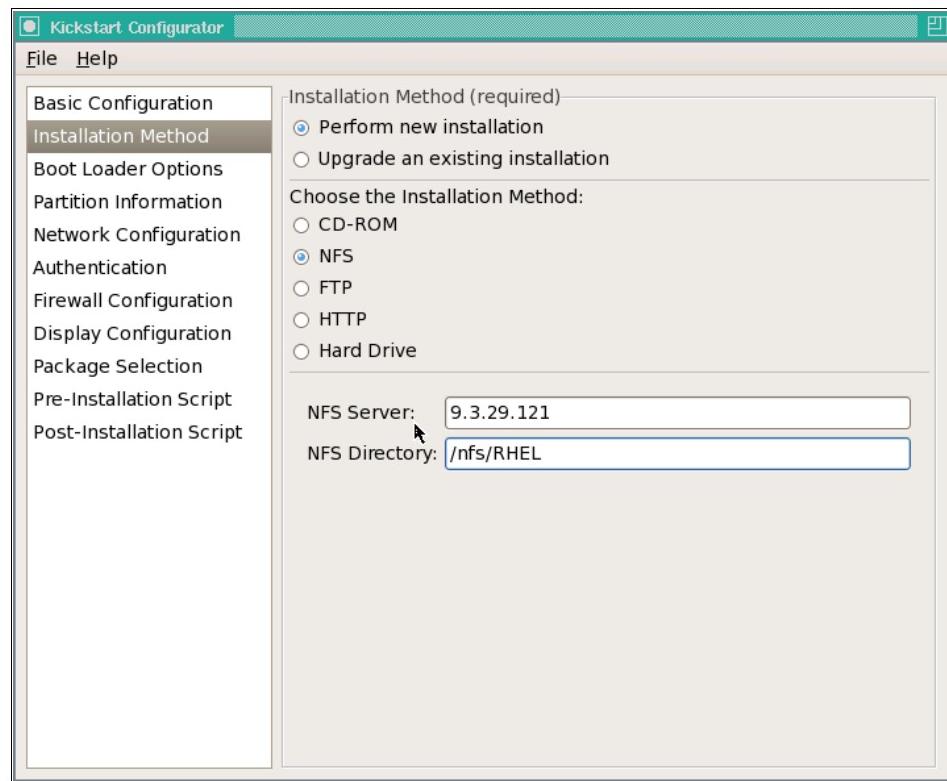


Figure 0-2 Installation Method panel (©2008 Red Hat, Inc.)

4. The next editable panel is the Partition Information panel, shown in Figure 0-3. Press **Add** to create a partition. The tool will help you select the mount point, file system type, and partition size.

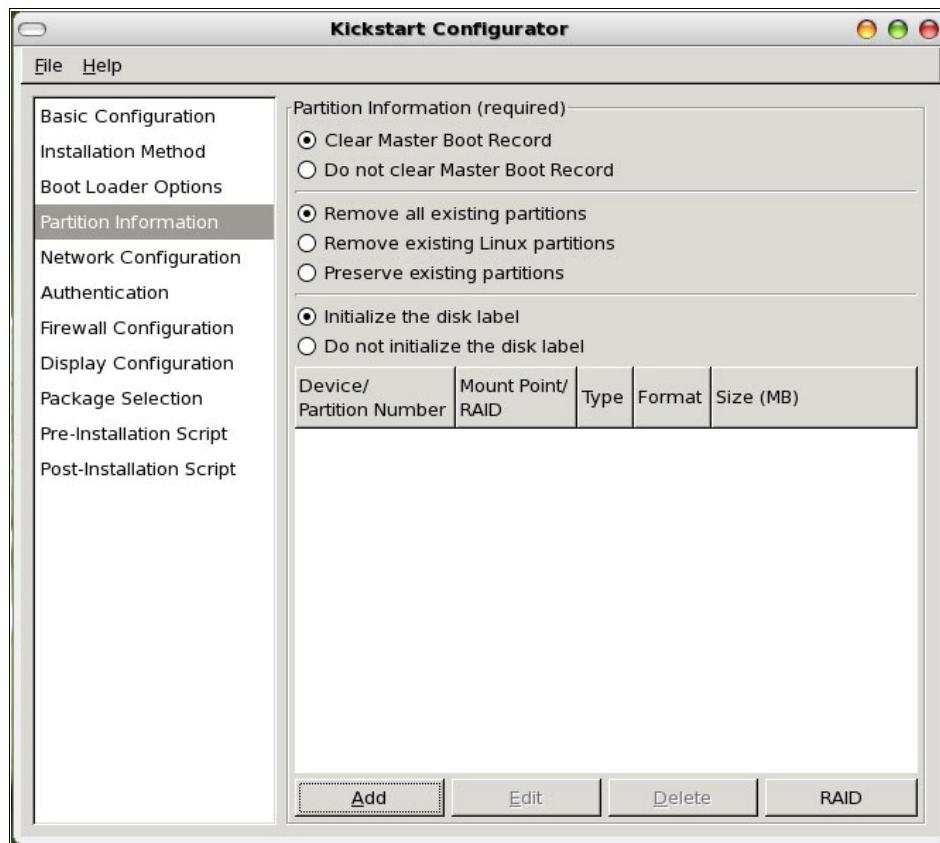


Figure 0-3 Partition Information panel (©2008 Red Hat, Inc.)

5. In the Network Configuration panel, click **Add Network Device** to add the devices you are installing from. If you need to go back and make changes to this setup, click **Edit Network Device** (see Figure 0-4).

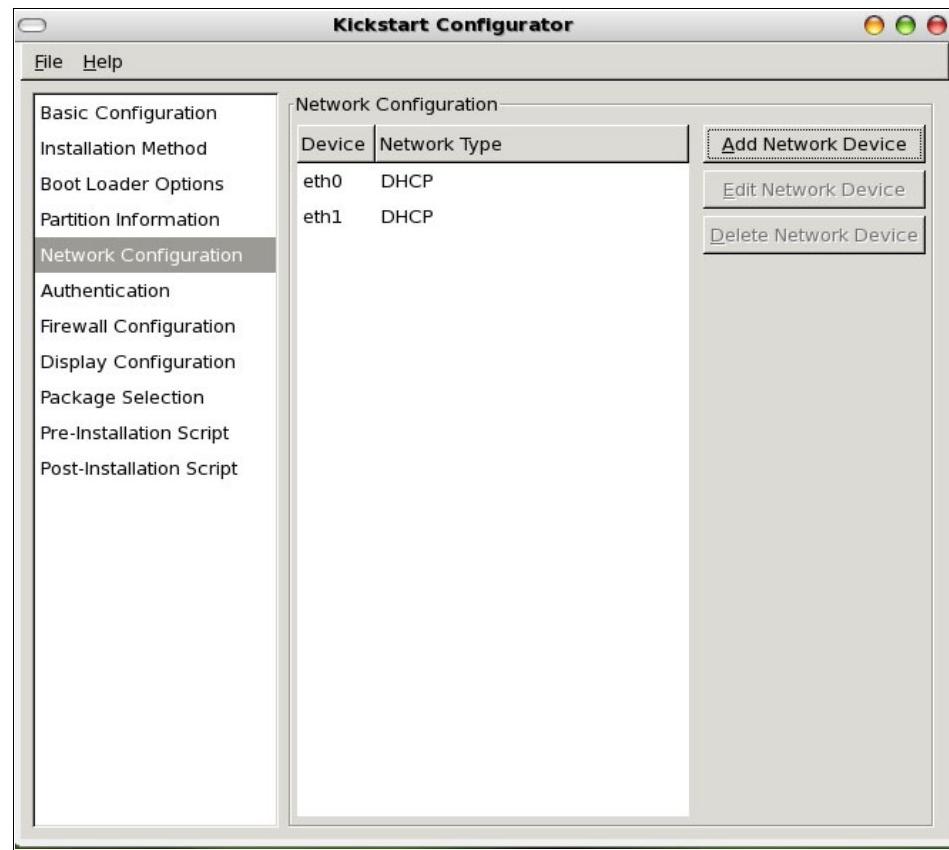


Figure 0-4 Kickstart Configurator Network Configuration panel (©2008 Red Hat, Inc.)

6. The next panel is the Authentication panel. In this configuration, we use the default settings.
7. Figure 0-5 shows the Firewall Configuration panel. As an example, it is good to enable SSH and to trust interface eth1 at the very minimum to access the system later using the network.

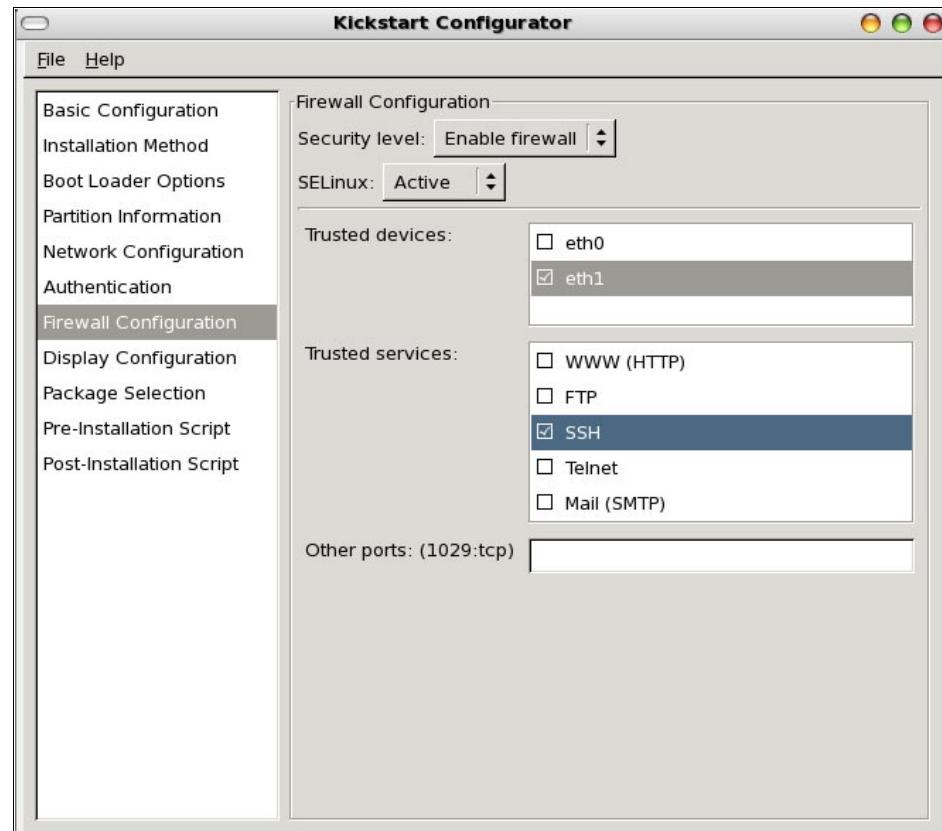


Figure 0-5 Firewall Configuration panel (©2008 Red Hat, Inc.)

8. Figure 0-6 shows the Package Selection panel. It is not possible to select individual packages from this panel. However, you can add individual packages to the %packages section of the Kickstart file after saving it.

Note: If you see the message “*Package selection is disabled due to problems downloading package information*” in the Package Selection panel, it means you have no repositories defined. You can either register the system with Red Hat Network (RHN) to a local internal RHN Satellite through the yum-rhn-plugin, or copy all the files from the Red Hat Enterprise Linux 5.3 installation discs to a directory and use them as a reference as given in the Red Hat Knowledge Base article available at:

http://kbase.redhat.com/faq/FAQ_103_12748.shtm

9. After the package selection, save the configuration using the **File → Save** option in the menu.



Figure 0-6 Package Selection panel (©2008 Red Hat, Inc.)

10. The basic Kickstart configuration file created with the Kickstart Configurator is shown in Example 0-1.

Important: The order of the main sections in the Kickstart configuration file is important for the functionality.

Example: 0-1 Basic Kickstart configuration file

```
#platform=IBM pSeries
# System authorization information
auth --useshadow --enablemd5
# System bootloader configuration
bootloader --location=mbr
# Clear the Master Boot Record
zerombr
# Partition clearing information
clearpart --all --initlabel
# Use text mode install
text
# Firewall configuration
firewall --enabled --trust=eth0,eth1
# Run the Setup Agent on first boot
firstboot --disable
# System keyboard
keyboard us
# System language
lang en_US
# Installation logging level
logging --level=info
# Use NFS installation media
nfs --server=9.3.29.121 --dir=/nsf/RHEL
# Network information
network --bootproto=dhcp --device=eth0 --onboot=on
network --bootproto=dhcp --device=eth1 --onboot=on
# Reboot after installation
reboot
#Root password
rootpw --iscrypted $1$ug.Uehug$WGBeIYgPLXt8A3QcsnDa90

# SELinux configuration
selinux --enforcing
# Do not configure the X Window System
skipx
# System timezone
timezone America/New_York
# Install OS instead of upgrade
install

%packages
@sound-and-video
@gnome-desktop
```

11. Manually adjust the Kickstart configuration file that you have created with a text editor if desired.

Note: If you have not defined any disk partition options or you were unsure of your disk partition layout, we recommend that you manually edit the Kickstart file to include the following information after the #Partition clearing information section:

```
#Disk partitioning information
autopart
```

This option will automatically create disk partitions.

Red Hat Enterprise Linux 5.3 Installation Guide is a good source of information regarding all Kickstart file options. It is also available at the following Web site:

http://www.redhat.com/docs/en-US/Red_Hat_Enterprise_Linux/5.3/html/Installation_Guide/ch-kickstart2.html

8.5.3 Performing a Kickstart installation

It is possible to run unattended installations with the kickstart file stored on a floppy disk, on a CD/DVD, on an USB key, on a partition in the system, or on some NFS/HTTP/FTP server. In this section we provide instructions on how to run an automatic installation with the kickstart file stored on a NFS exported directory, accessible through the network.

Pass the location of the Kickstart file during boot. BOOTP does not have the ability to provide anything more than the location to the bootable image and the server IP address hosting the Kickstart file. Use the Open Firmware prompt to pass the required parameters:

1. Type 8 during the LPAR boot process to go to the Open Firmware prompt, as shown in Figure 8-23.

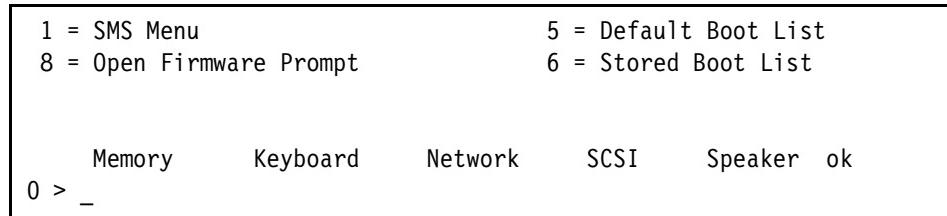


Figure 8-23 Open Firmware prompt

2. Type the following command in the Open Firmware prompt to start automated installation. For example, if the configuration file is served using NFS:

```
boot net ks=nfs://192.168.1.254/ks.cfg ksdevice=eth1 ip=dhcp
```

Press the Enter key and the process will begin. The automated Red Hat Enterprise Linux installation is now complete.

Tip: *Red Hat Enterprise Linux 5.3 Installation Guide* is a good source of information regarding all Kickstart file options. It is available at:

http://www.redhat.com/docs/en-US/Red_Hat_Enterprise_Linux/5.3/html/Installation_Guide/ch-kickstart2.html

8.6 IBM service and productivity tools

The IBM service and productivity tools are packages that are installed after a successful Red Hat Enterprise Linux installation.

Important: These packages are *not* shipped with the Red Hat installation CDs. They are IBM owned and distributed.

These packages enable features such as:

- Reliability, availability, and serviceability (RAS) functionality
- I/O hotplug
- Dynamic Logical Partitioning (DLPAR) capabilities
- Live partition migration capabilities

See Appendix D, “Service and productivity tools for Linux” on page 545 for more information on installing the service and productivity tools specific to your system’s configuration.



SUSE Linux Enterprise Server V11 installation

This chapter describes the procedures to install SUSE Linux Enterprise Server (SLES) V11 on a JS43 BladeCenter.

We discuss the following topics:

- ▶ “Supported operating systems” on page 368
- ▶ “Considerations and prerequisites” on page 368
- ▶ “Linux LPAR installation using DVD” on page 369
- ▶ “Linux network installation (detailed)” on page 374
- ▶ “Native SLES 11 installation” on page 394
- ▶ “SLES 11 automated installation” on page 395
- ▶ “IBM service and productivity tools” on page 396

9.1 Supported operating systems

SUSE Linux Enterprise Server 10 Service Pack 1 (SLES 10 SP1) for POWER or later supports installation on a JS43.

This chapter specifically covers installing SUSE Linux Enterprise Server 11 for POWER with a DVD and over the network on a PowerVM LPAR.

9.2 Considerations and prerequisites

There are some system configuration considerations and prerequisites prior to installing SLES 11 on a JS43 partition. They are covered here.

9.2.1 PowerVM LPAR considerations

By the characteristics of the virtualization features in the System p JS43 Blade Center, the operating system and applications do not know they are running in a micro-partitioned or virtualized I/O environment. This allows applications to run unmodified in a partition that takes advantage of virtualization features.

Because the PowerVM partition handles the transition of the virtual adapter's I/O operation to the physical adapter, it is important to guarantee that the partition is properly sized to handle the I/O requirements. A good source for processor and memory requirements for PowerVM partitions based on I/O requirements is found at:

<http://www14.software.ibm.com/webapp/set2/sas/f/vios/documentation/perf.html>

Additional information about virtualization and LPAR setup is available in Chapter 4, "System planning and configuration using VIOS with IVM" on page 71.

9.2.2 SUSE Linux Enterprise Server 11 installation considerations

The SUSE Linux Enterprise Server 11 document recommends at least 512 MB of RAM to install SUSE Linux Enterprise Server 11.

Tip: We recommend at least 1-2 GB per PowerVM LPAR and 4 GB on LPARs running CPU intensive loads.

In addition, the SLES 11 installation guide suggests to have at least 1.5 GB of hard disk space or have one or more hard disk partitions that can be deleted to free up the miNFSum disk space for the Linux installation.

Tip: We recommend 10 GB or more *total* hard disk space for each PowerVM LPAR.

The Novell Web site has additional installation preparation information for SLES 11 available at:

<http://www.novell.com/documentation/sles11/index.html>

9.3 Linux LPAR installation using DVD

With the PowerVM installed and the system partitioned into LPARs using the PowerVM LPAR considerations and the SUSE Linux Enterprise Server 11 prerequisites, we are ready to install Linux on the JS43 Blade Center LPAR.

Note: Even though this section covers installation via the Integrated Virtualization Manager (IVM) console, there are other console options available on the JS43; these are covered in Appendix A, “Consoles, SMS, and Open Firmware” on page 493.

Start an LPAR installation of SLES 11 using a DVD

The following directions are for installing SLES 11 from a DVD:

1. On the Advanced Management Module (AMM) select **Blade Tasks** → **Remote Control**.
2. In the right panel look for Media Tray Owner and click the drop-down arrow to display the different blade bays. Select your blade bay location with the mouse as shown in Figure 9-1.

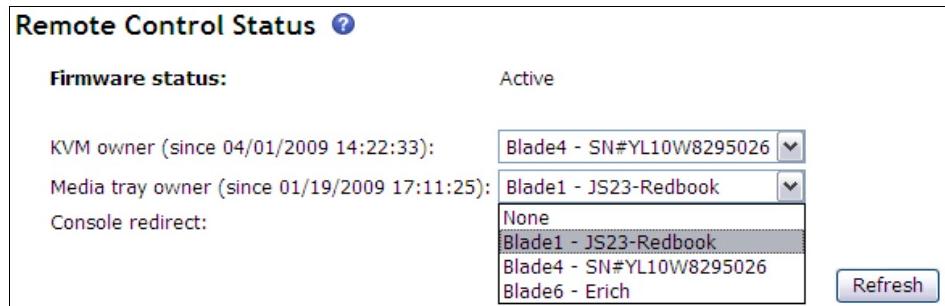


Figure 9-1 Start Remote Console panel

3. Press Refresh.

Note: The other option is to press the **MT** button on the blade to assign the media tray to the blade.

Important: Make sure no other blade in the Blade Center is using the media tray before pressing this button. The blade's MT light is on if the media tray is assigned to them.

4. Double-check that your blade bay owns the media tray by opening the AMM panel and selecting **Monitors** → **System Status**. The left side will show an X in the MT column of your blade bay location. Figure 9-2 gives one example of this assignment.

Blades ?							
Click the icon in the Status column to view detailed information about each blade.							
Bay	Status	Name	Pwr	Owner **		cKVM *	I/O Compatibility
				KVM	MT *		
1	■	JS23-Redbook	On	✓	✓	OK	

Figure 9-2 The checkmark indicates that the Media Tray (MT) is owned by the blade in Bay 1

5. Place the DVD into the Blade Center media tray.
6. Log in to the Integrated Virtualization Manager (IVM) using a Web browser.
 - a. Select one of the available PowerVM LPARs to install with SLES 11 by placing a checkmark in the Select box.
 - b. Click **Activate** as shown in Figure 9-3

Partition Details										
	ID	Name	State	Uptime	Memory	Processors	Entitled Processing Units	Utilized Processing Units	Reference Code	
<input type="checkbox"/>	1	is23-vios	Running	20.33 Hours	1 GB	4	0.4	0.57		
<input checked="" type="checkbox"/>	2	is23-lp1	Not Activated		1 GB	1	0.1		00000000	

Figure 9-3 Activating an IVM partition

- c. Click the drop-down arrow to the right of the More Tasks field and select **Open terminal window**.

Important: Make sure the latest Java Runtime Environment (JRE) is installed on the native system to run the IVM terminal. At the time of this publication, the recommended JRE is Sun's JRE 1.6.0_13, or higher.

Partition Details										
	ID	Name	State	Uptime	Memory	Processors	Entitled Processing Units	Utilized Processing Units	Reference Code	
<input type="checkbox"/>	1	is23-vios	Running	20.33 Hours	1 GB	4				
<input checked="" type="checkbox"/>	2	is23-lp1	Not Activated		1 GB	1				

Figure 9-4 Opening a terminal window from the IVM

The console is a pop-up and it will ask you to authenticate with the PowerVM User ID and password.

7. The System Maintenance Services (SMS) menu appears in the IVM terminal.

Note: The SMS menu appears very quickly after activating the box, so have your hand ready on the 1 key so you can press it immediately when you see the SMS menu. If you happen to miss the SMS option section window, just shut down and reactivate the LPAR again. The IVM terminal will hold the connection to the PowerVM LPAR even if the LPAR is shut down.

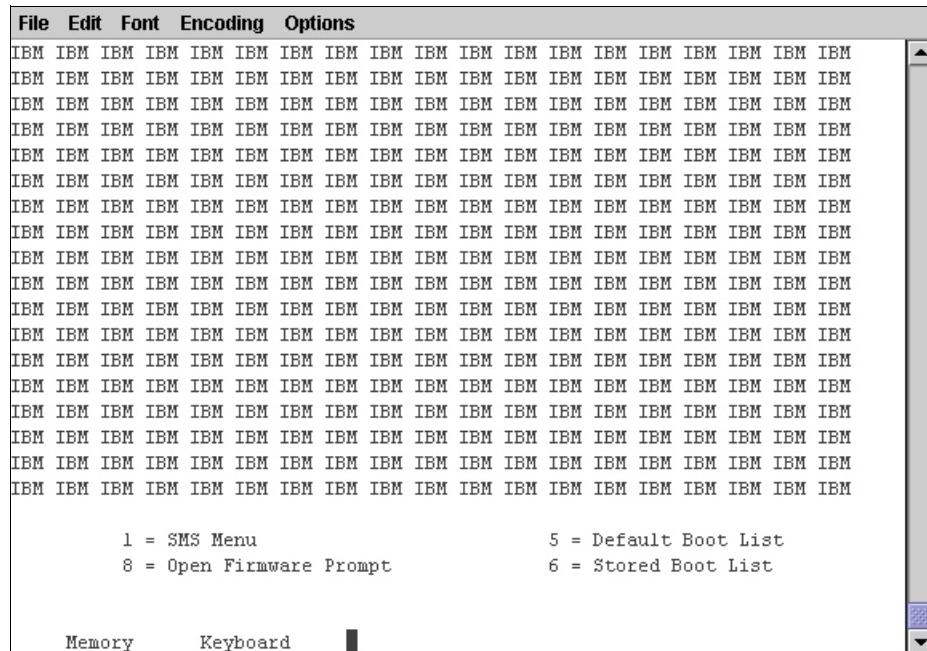


Figure 9-5 SMS Menu

- a. Select **1 = SMS Menu** by pressing the number 1 on the keyboard.

Note: Press the number next to the desired system function to navigate through the SMS menu.

- b. Select option **5. Select Boot Options**.
- c. Choose option **1. Select Install/Boot Device**.
- d. Pick **3. CD/DVD**.
- e. Select **6. USB**.
- f. Finally, select **1. USB CD-ROM**. See Figure 9-6 on page 372.

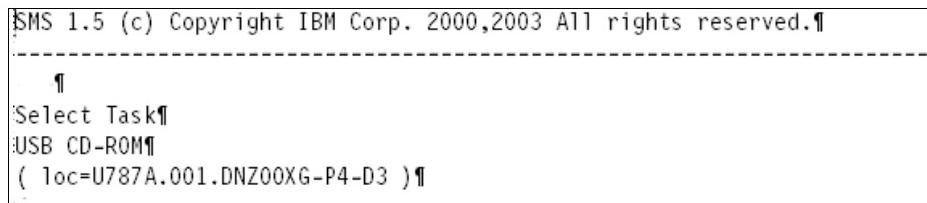


Figure 9-6 Select CD-ROM

g. Choose **2. Normal Mode Boot**. See Figure 9-7 on page 373.

```
1. Information
2. Normal Mode Boot
3. Service Mode Boot

-----
| Navigation keys:
M = return to Main Menu
ESC key = return to previous screen X = eXit System Management Services
-----
| Type the number of the menu item and press Enter or select Navigation
Key:|
```

Figure 9-7 Select Mode Boot

- h. Pick **1.Yes** to exit the SMS menu.
- i. At the Linux **boot:** prompt, type `install`, See Figure 9-8 on page 373 then press Enter to confirm. The LPAR will start reading from the DVD, which can take a couple of minutes.

```
Welcome to SuSE:SLE-11:GA!
Type "install" to start the YaST installer on this CD/DVD
Type "slp" to start the YaST install via network
Type "rescue" to start the rescue system on this CD/DVD
Welcome to yaboot version 1.3.11.SuSE
Enter "help" to get some basic usage information
boot:|
```

Figure 9-8 Select installation type

8. After successfully booting off the DVD, the SUSE language selection panel appears on the IVM terminal and Your awesome Setup Tool (YaST) will assist with the completion of the install. More detailed installation instructions are available here:

http://www.novell.com/documentation/sles11/book_sle_admin/index.html?page=/documentation/sles11/book_sle_admin/data/book_sle_admin_pre.html

9.4 Linux network installation (detailed)

This section describes a Network File System (NFS) installation on a PowerVM LPAR using an external Storage Area Network (SAN) device. Even though this example uses a specific storage device and network option to complete the installation, this network installation process is generic enough that it will work for all supported internal or external storage devices on JS43, and only a few slight modifications are needed in the YaST installer setup if a different network installation option is selected. The goal of this section is to show the SAN disk install in detail so the steps can be referenced later.

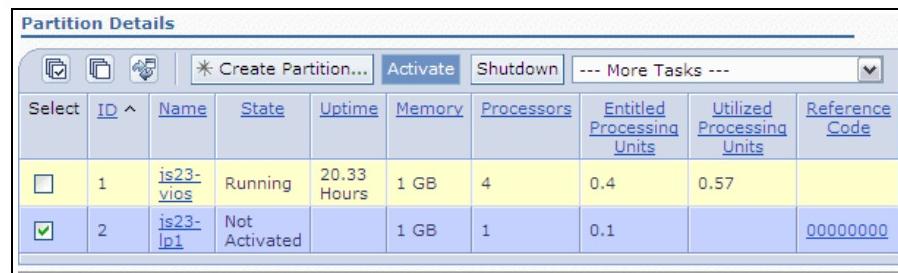
This installation is using an NFS installation method, but NFS is one of many supported network installation types on Linux. The NFS server configuration and setup are described in detail in , “Installing Linux using the network - General remarks” on page 536 of Appendix C, “Additional Linux installation configuration options” on page 535.

This section assumes that the user has already set up the NFS server properly and has read 9.2.1, “PowerVM LPAR considerations” on page 368 and followed the installation prerequisites in 9.2.2, “SUSE Linux Enterprise Server 11 installation considerations” on page 368.

Start an LPAR installation of SLES 11 over the network

The following instructions guide you through the installation of SLES 11 through a network.

1. Log in to the Integrated Virtualization Manager (IVM) using a Web browser.
 - a. Select one of the available PowerVM LPARs to install with SLES 11 by placing a checkmark in the Select box.
 - b. Click **Activate** as shown in Figure 9-9 on page 374.



Partition Details										
Select	ID ^	Name	State	Uptime	Memory	Processors	Entitled Processing Units	Utilized Processing Units	Reference Code	More Tasks
<input type="checkbox"/>	1	is23-vios	Running	20.33 Hours	1 GB	4	0.4	0.57		
<input checked="" type="checkbox"/>	2	is23-lp1	Not Activated		1 GB	1	0.1		00000000	

Figure 9-9 Activating an IVM partition

- c. Click the drop-down arrow to the right of the More Tasks field and select **Open terminal window** as shown in Figure 9-10.

Important: Make sure the latest Java Runtime Environment (JRE) is installed on the native system to run the IVM terminal. At the time of this publication, the recommended JRE is Sun's JRE 1.6.0_13, or higher.

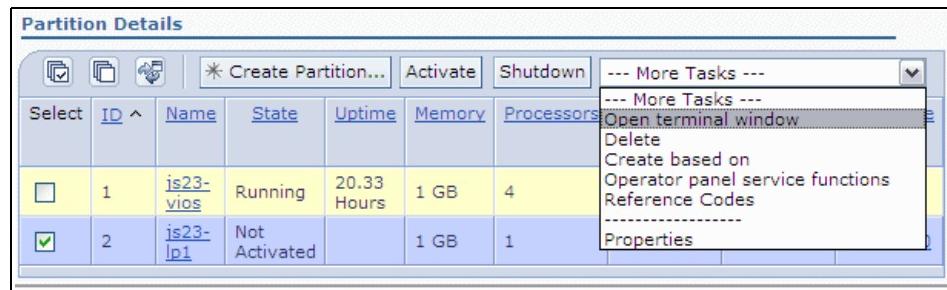


Figure 9-10 Opening a terminal window from the IVM

The console is a pop-up and it will ask you to authenticate with the PowerVM User ID and password.

2. The firmware boot panel appears in the IVM terminal.

Note: The SMS menu option appears very quickly after activating the LPAR, so have your hand ready on the 1 key so you can press it immediately when you see the SMS menu. If you happen to miss the SMS selection window, just shut down and reactivate the LPAR again. The IVM terminal will hold the connection to the PowerVM LPAR even if the LPAR is shutdown.

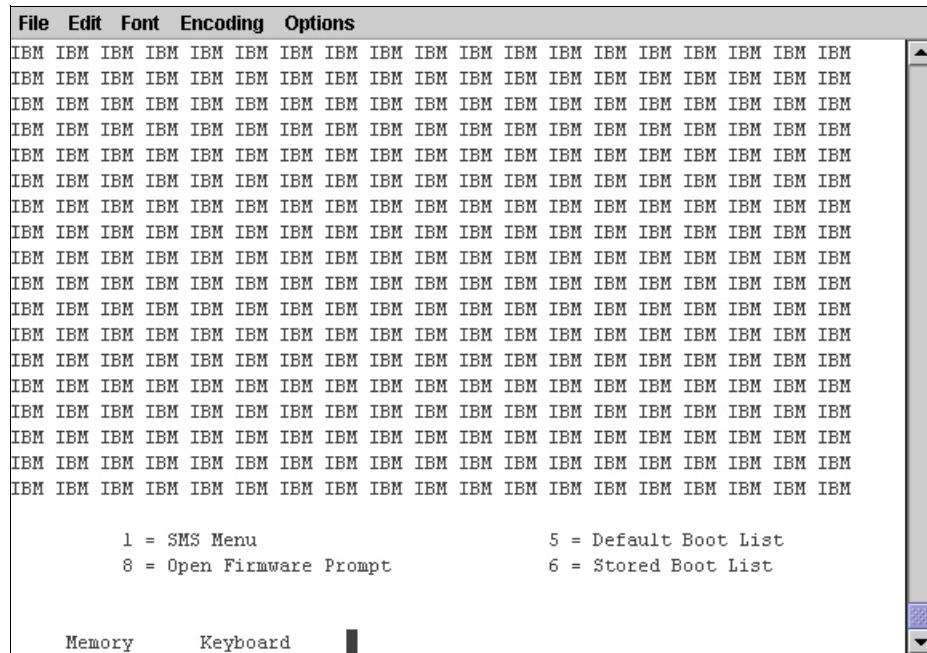


Figure 9-11 SMS Menu

- Select **1 = SMS Menu** by pressing the number 1 on the keyboard.

Note: Press the number next to the desired system function to select and navigate through the SMS menu.

- Choose **5. Select Boot Options**.
- Choose **1. Select Install/Boot Device**.
- Choose **6. Network**.
- Pick **1. BOOTP**.
- Choose a network port.
- Choose **2. Normal Mode Boot**.
- Pick **1.Yes** to exit the SMS menu.

If everything is set up correctly, the system connects using DHCP to the TFTP service to start loading the boot image for the first stage of the install. In a couple of minutes the Main Menu panel appears, as shown in Figure 9-12 on page 377.

- Select **3) Expert**, as shown in Figure 9-12 on page 377.

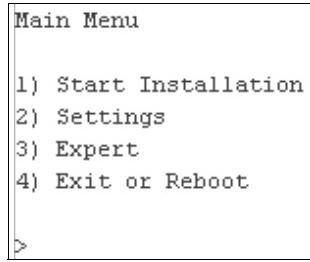


Figure 9-12 Main Menu

Tip: Press the number next to the desired configuration option and then the Enter key to select it in the Main Menu window.

The Enter key alone will move you back to the previous option window.

2. Choose **2) Kernel Modules (Hardware Drives)**, as shown in Figure 9-13 on page 377.

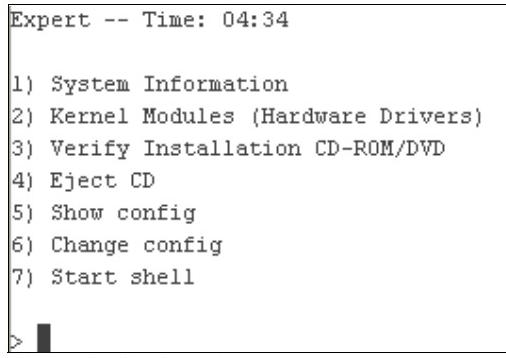


Figure 9-13 Expert

3. Choose **1) Load ppc Modules**, as shown in Figure 9-14 on page 378.

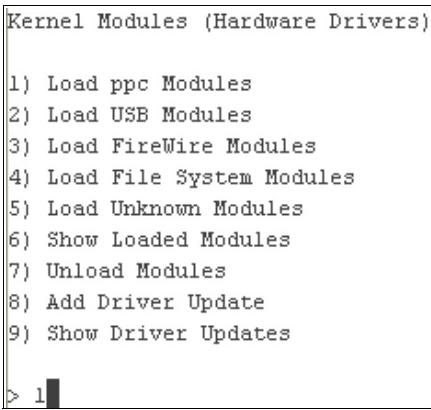


Figure 9-14 Load ppc Modules

4. Select each individual module to pre-install based on your LPAR's network configuration. Press the number next to the module name and then the Enter key, then press the Enter key again to confirm.

Tip: Use the up/down scroll bar on the IVM terminal to navigate the module list. The most commonly used modules are 5) e1000 : Intel PRO/1000, 15) ehea : EHEA and IBMVETH.

5. Press the **Enter** key after you have finished loading the modules to go back to the main menu.
6. Choose **1) Start Installation or Update**, as shown in Figure 9-15.

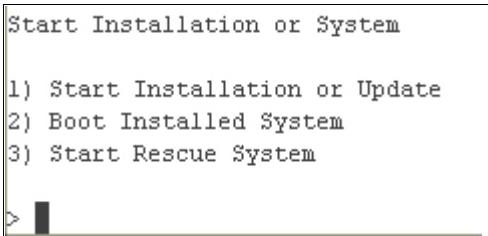


Figure 9-15 Start installation or update option

7. Select **2) Network**.
8. Select **3) NFS** as the network protocol.
9. Choose **1) Yes** or **2) No** for network configuration via DHCP.
 - a. If option **2) No** was selected, enter the static IP information:

- LPAR's IP address
- LPAR's netmask
- LPAR's gateway
- LPAR's name server
- The NFS server's IP addressThe directory on the NFS server which contains the SLES 11 image. Figure 7-16 shows a sample configuration.

```

Choose the network protocol.

1) FTP
2) HTTP
3) NFS
4) SMB / CIFS (Windows Share)
5) TFTP

> 3
Detecting and loading network drivers

Automatic configuration via DHCP?

1) Yes
2) No

> 2

Enter your IPv4 address.
Example: 192.168.5.77/24
> 9.8.234.221

Enter your netmask. For a normal class C network, this is usually
255.255.255.0
[255.255.255.0]> 255.255.255.128

Enter the IP address of the gateway. Leave empty if you don't need one
> 9.8.234.138

Enter your search domains, separated by a space:
> 9.8.234.180

Enter the IP address of your name server. Leave empty or enter "+++" if you
don't need one
> 9.8.234.180

Enter the IP address of the NFS server
> 9.8.234.180

Enter the directory on the server
[/]> /srv/repository/suse/install/sles-11-ppc

```

Figure 9-16 Static network configuration example

The LPAR begins reading from the SLES 11 image directory and then displays the Your awesome Setup Tool (YaST) Welcome panel, as shown in Figure 9-17 on page 380.

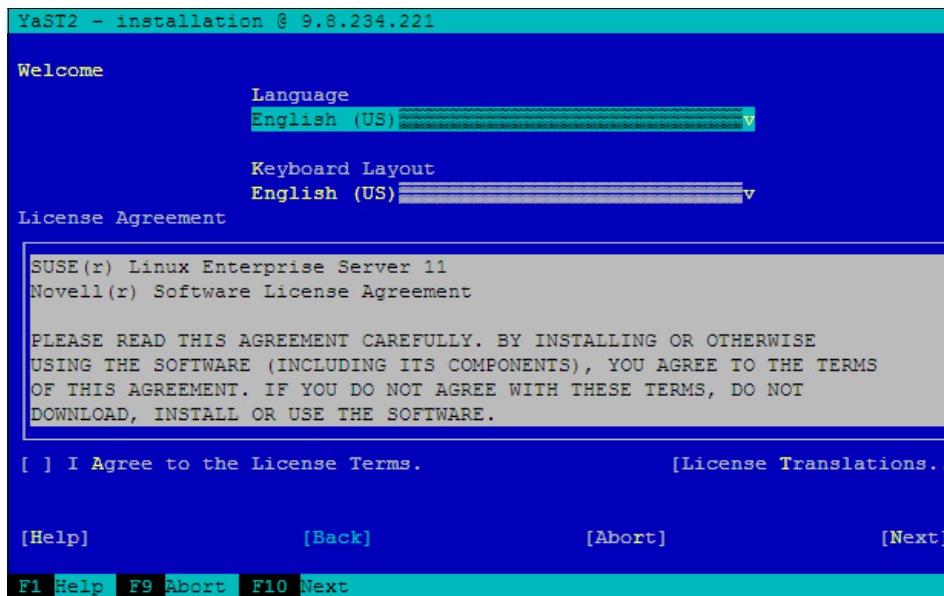


Figure 9-17 YaST Welcome panel

Tip: Navigate the YaST tool by using the Tab key to move between sections, the up/down arrow keys to move within a specific window section, the space bar to check a "()" entry with an "x," the Enter key to confirm a selection with square brackets "[]" around it, and the Delete key to erase entries.

10. Highlight **English** in the Language selection panel. Press the Tab key to **[Next]** and then press Enter.
11. On the next panel, select **Yes, I Agree to the Licence Agreement** and then press the Tab key to **[Next]** and press Enter to confirm.
12. Select **New Installation** for the Installation Mode, as shown in Figure 9-18 on page 381.

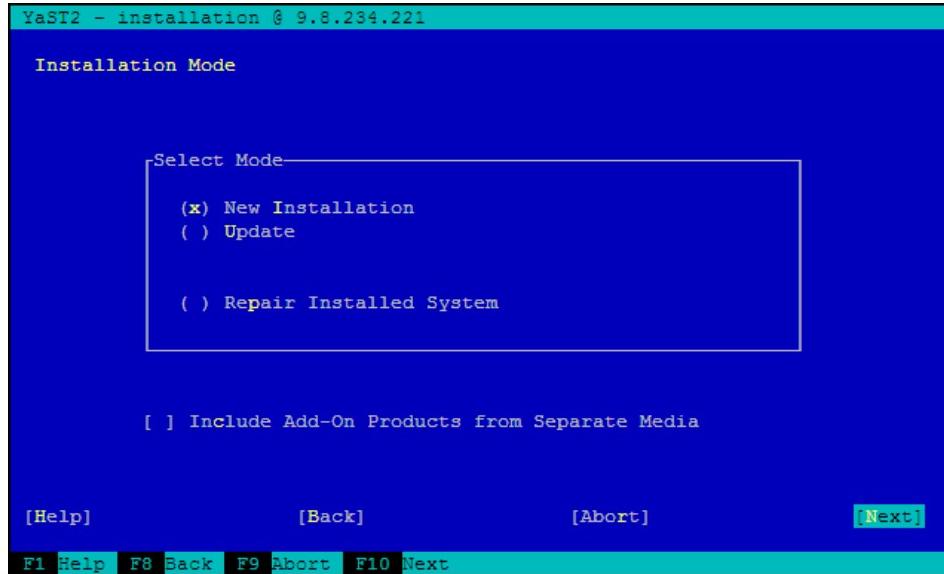


Figure 9-18 Installation Mode

13. Configure your clock and time zone information, as shown in Figure 9-19 on page 381.



Figure 9-19 Clock and Time Zone

14. The Installation Settings window provides the Keyboard layout, Partitioning information, Software installation options, and the install Language configuration. Select the [Change...] option to edit any of these fields. Select [Accept] when these settings are complete, as shown in Figure 9-20 on page 382.

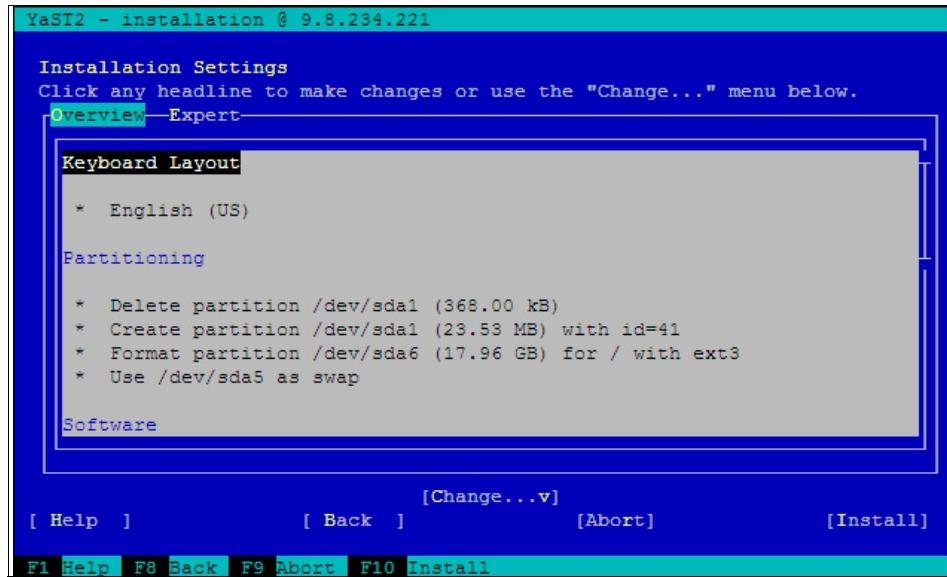


Figure 9-20 Installation Settings

15. Select [I Agree] to the AGFA Monotype Corporation License Agreement, as shown in Figure 9-21 on page 383.

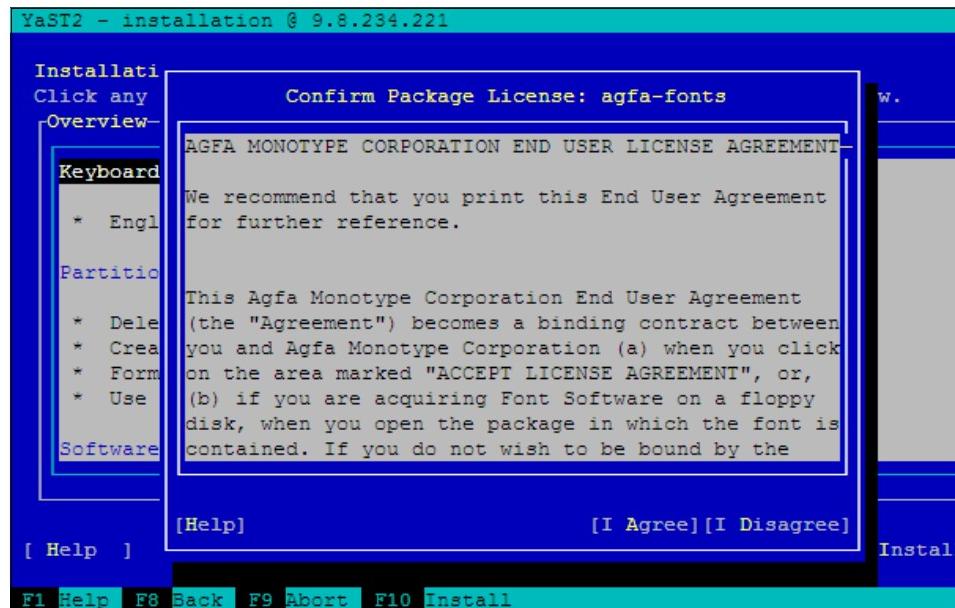


Figure 9-21 AGFA License Agreement

16. Choose [Install] to start the installation, as shown in Figure 9-22 on page 383.

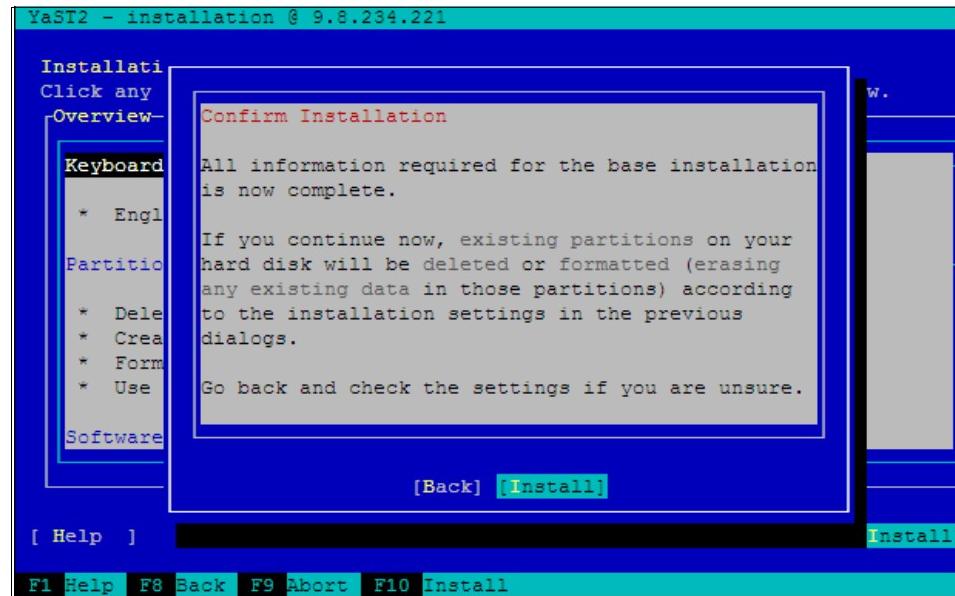


Figure 9-22 Confirm Installation

The YaST window refreshes to the installation progress bars, as shown in Figure 9-23. The top status bar shows the progress YaST has made installing a specific package and the bottom is the progress of the entire installation. The system will reboot after the installation completes.

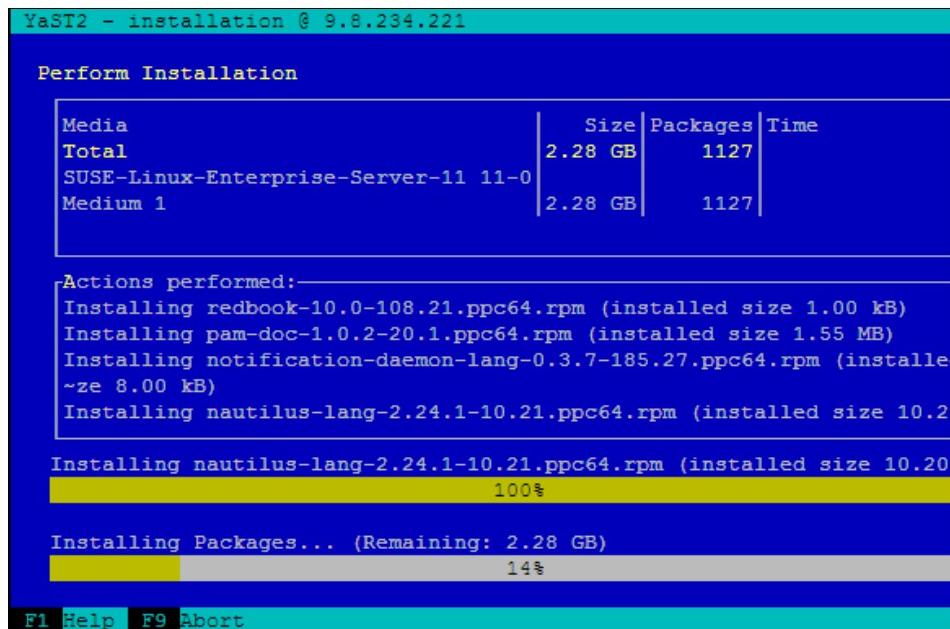


Figure 9-23 YaST installation progress window

Note: If the LPAR does not automatically boot from the intended hard disk (boot device) after reboot, try this:

- ▶ Shut down and reactivate the LPAR from the IVM.
- ▶ Enter the SMS Menu.
- ▶ Select **5. Select Boot Options → 1. Select Install/Boot Device → 5. Hard Drive → 9. List All Devices.**
- ▶ Choose the appropriate hard disk with the Linux image from the given list.
- ▶ Select **2. Normal Mode Boot → 1. Yes.**

17. After booting from the appropriate storage device YaST will start again. Select **[Continue]** for each hardware device YaST detects, as shown in Figure 9-24 on page 385.

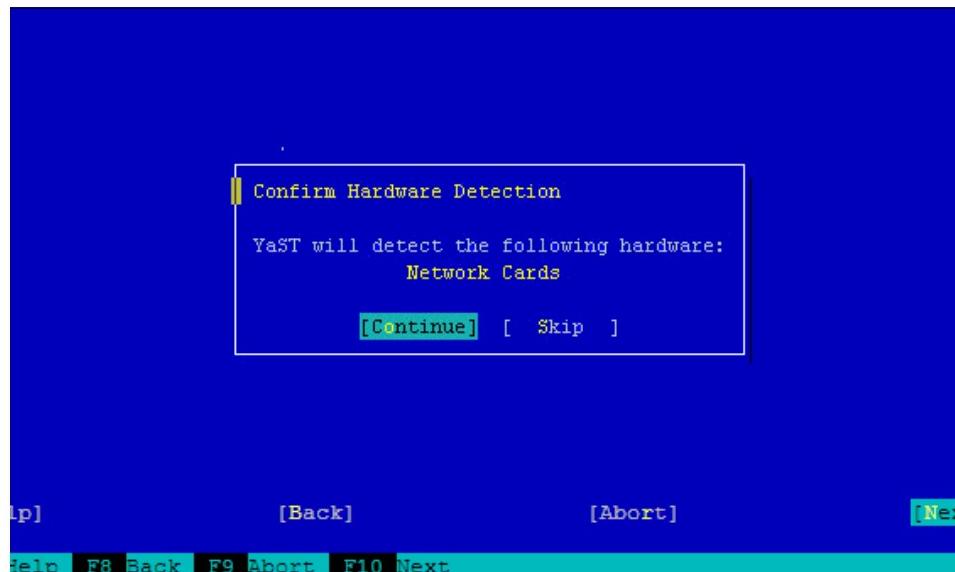


Figure 9-24 Confirm hardware detection window

18. Boot the system. See Figure 9-25 on page 385.

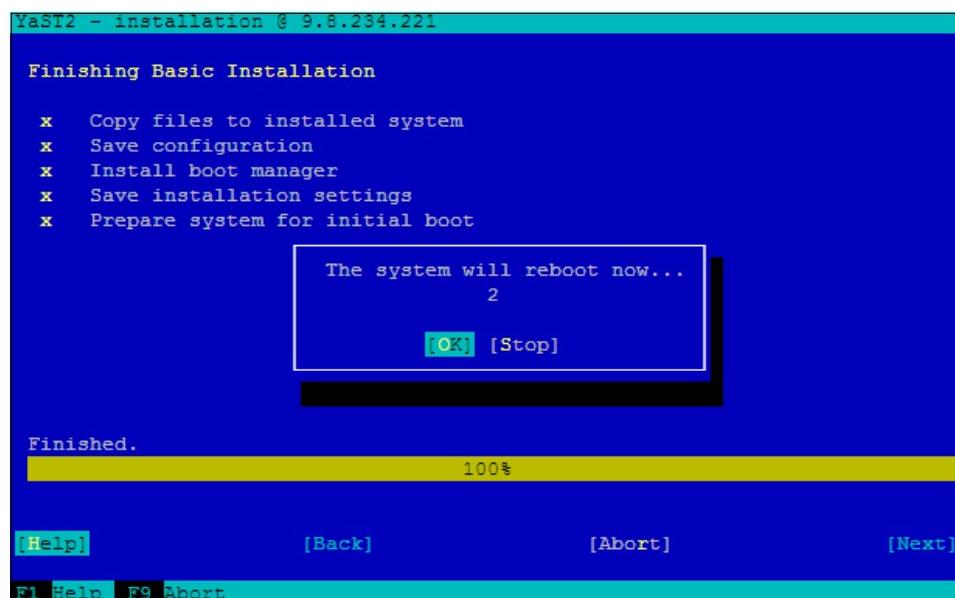


Figure 9-25 Reboot now

19. Enter the root user's password. Press [Next] to confirm, as shown in Figure 9-26 on page 386.

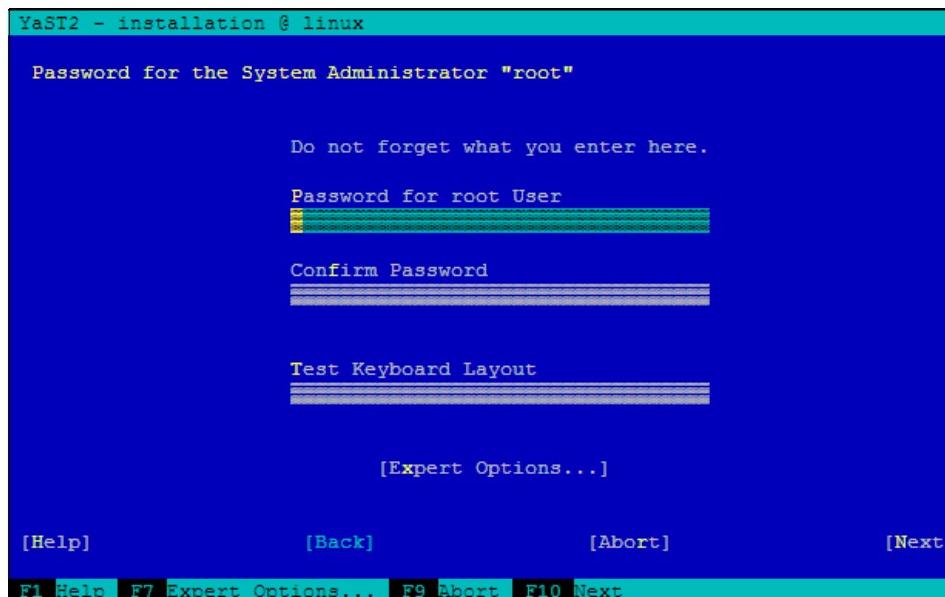


Figure 9-26 root User Password

20. Provide the hostname and the domain. Press [Next] to confirm. See Figure 9-27 on page 387.

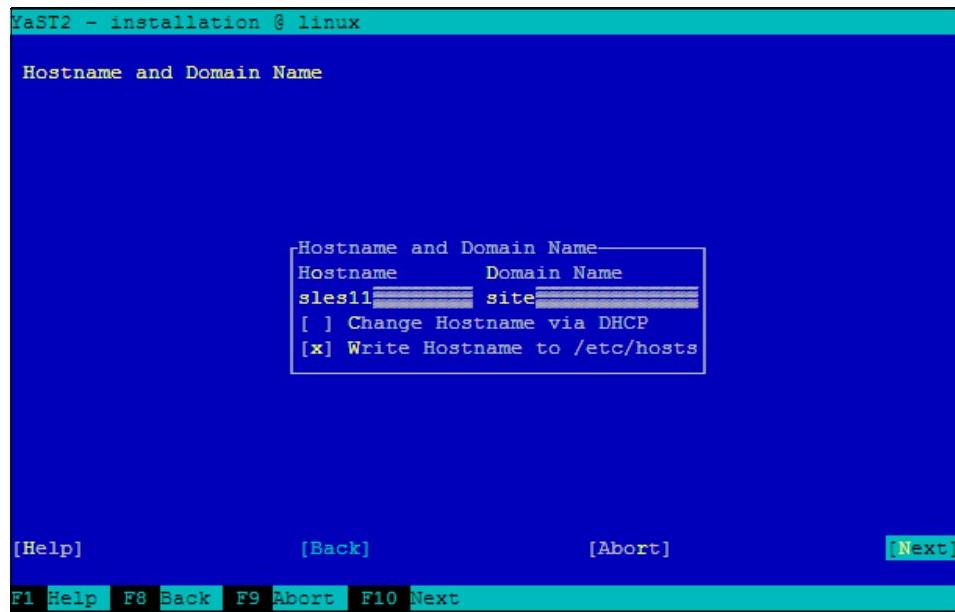


Figure 9-27 Hostname and Domain Name

21. Select **Use Following Configuration** in the Network Configuration window (Figure 9-28 on page 388) and verify that the Firewall is marked as **enabled**. Press the Tab key to **[Change....]** to change the Secure Shell (SSH) port settings to open.

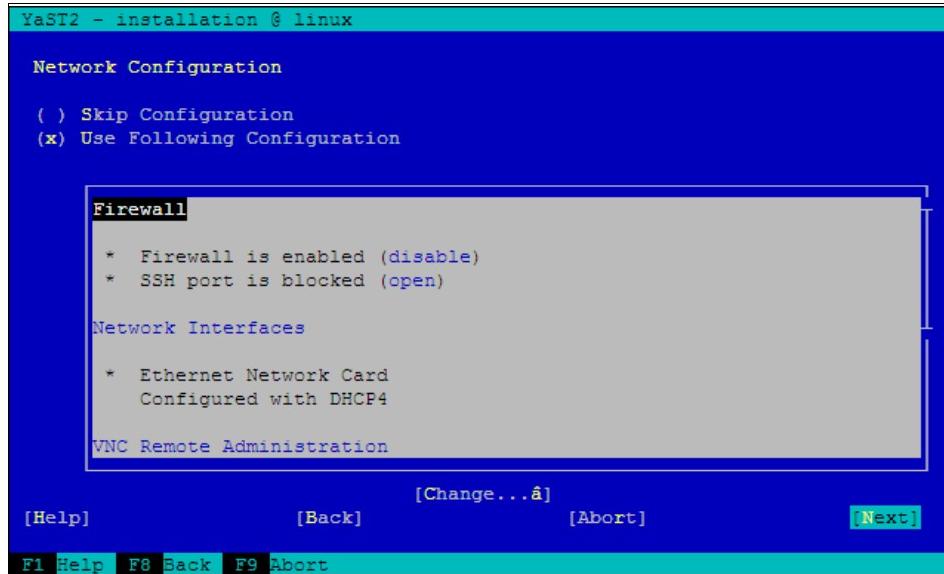


Figure 9-28 Change network configuration

- Select **Firewall** as shown in Figure 9-29.

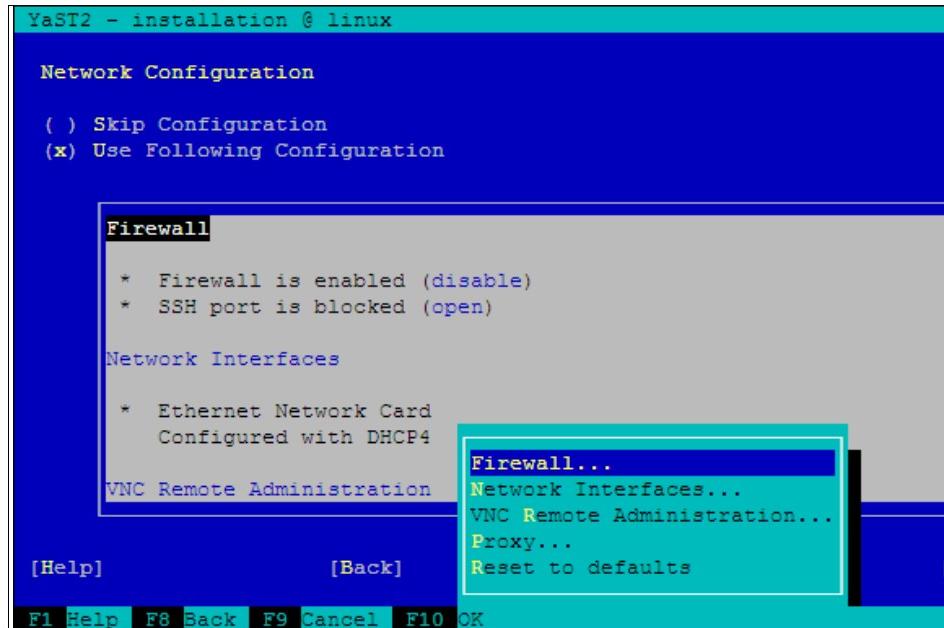


Figure 9-29 Select firewall

- b. Scroll to **Allowed Services**.
- c. Find and highlight **SSH** in the new window, as shown in Figure 9-30 on page 389. Finally, press Enter to confirm.

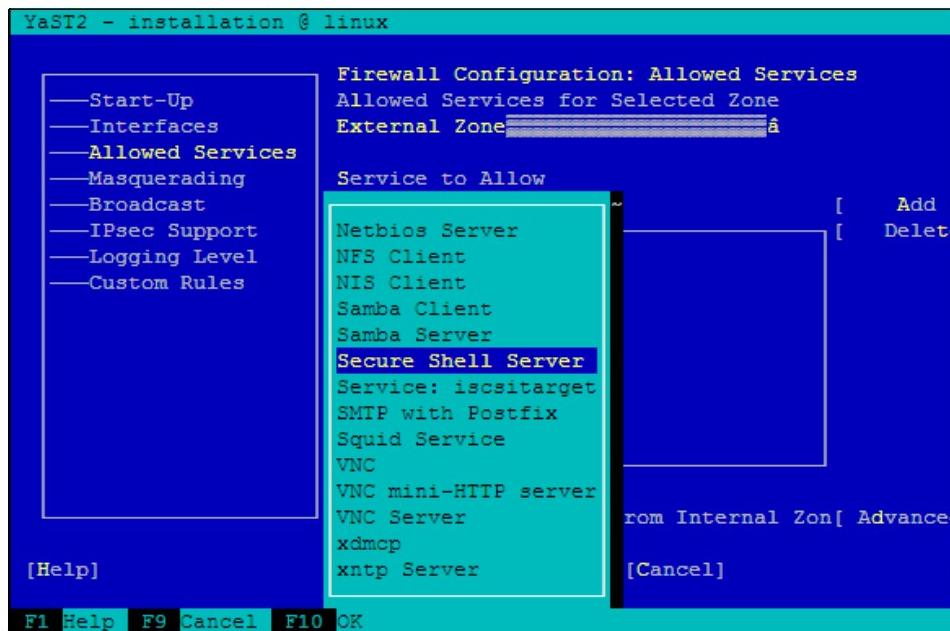


Figure 9-30 Services to allow list and selecting SSH service

- d. Press the Tab key to highlight [**Add**] and the press Enter to confirm.
- e. SSH will appear in the **Allowed Service** list, as shown in Figure 9-31 on page 390. Press [**Next**] to confirm.

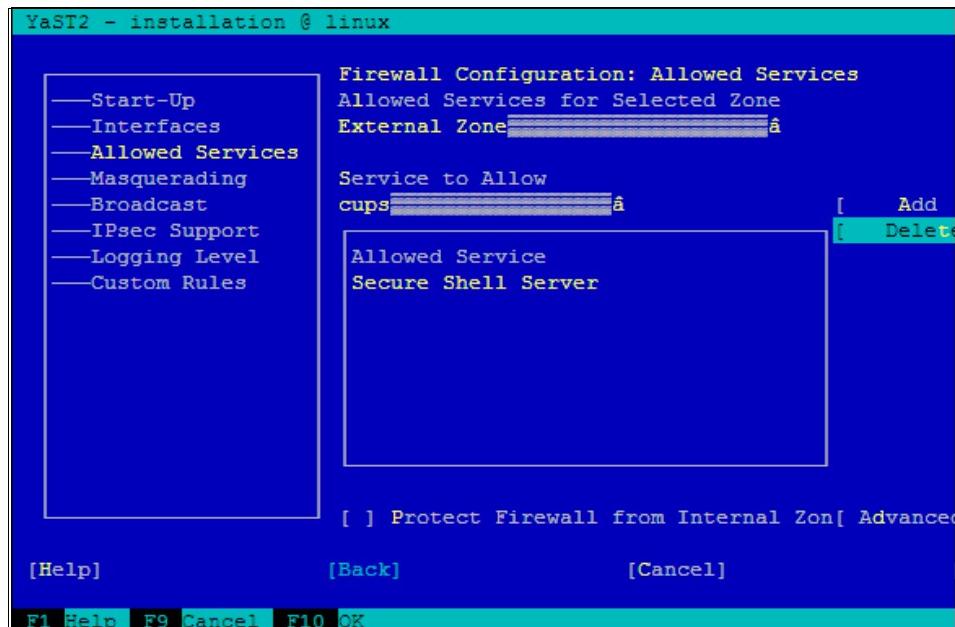


Figure 9-31 Allowed Service Secure Shell Server (SSH)

- f. Now the Firewall section of the Network Configuration window (Figure 9-32) shows “SSH port is open.”

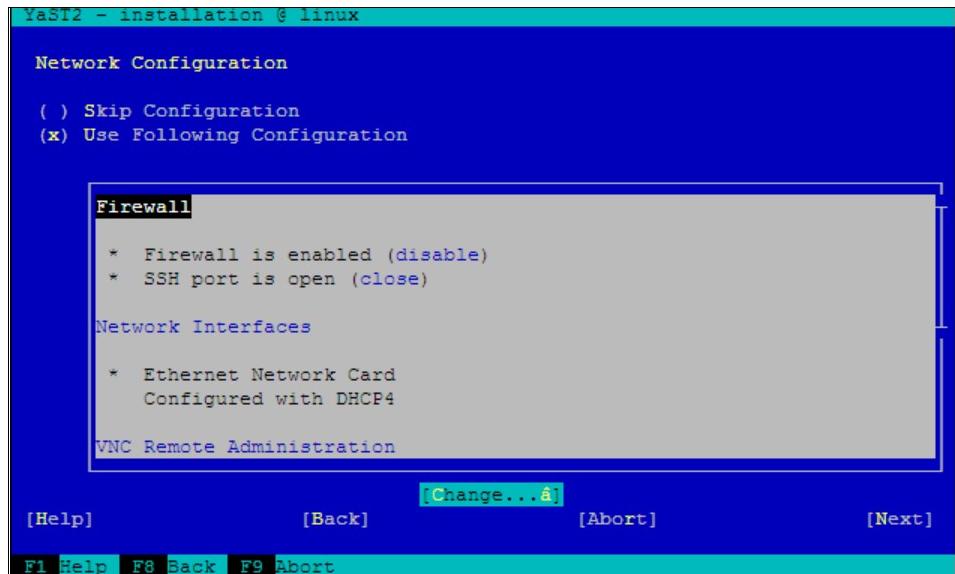


Figure 9-32 SSH port is open

22. Test the Internet connection, if desired.
23. Change the Certification Authority (CA) Installation setting, if desired. Select **[Next]** to confirm the changes.
24. Select the user authentication method appropriate for this LPAR and select **[Next]**. See Figure 9-33 on page 391.

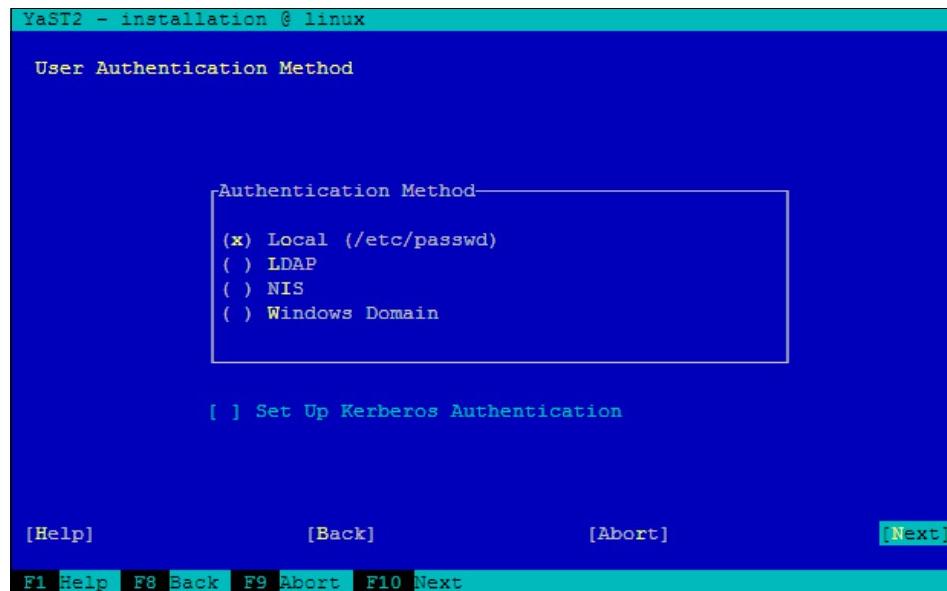


Figure 9-33 User Authentication Method

25. Create a local user and select **[Next]**. See Figure 9-34 on page 392.

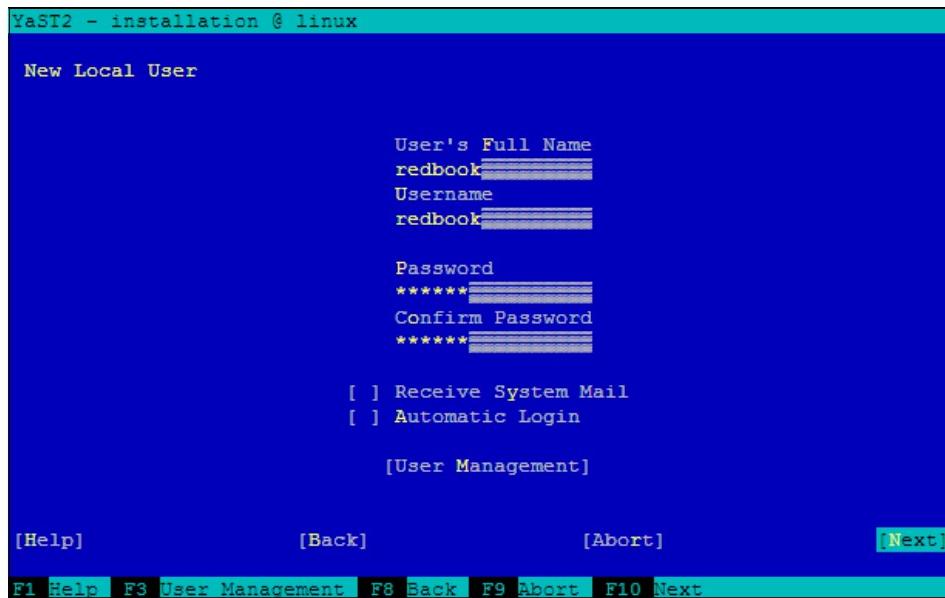


Figure 9-34 New Local User

- 26.YaST will write the configuration settings and then display the Release Notes.
Choose **[Next]** after reading the release notes.
- 27.Configure Hardware (Printers) if desired, then confirm the described configuration with **[Next]**.
- 28.YaST displays the Installation Completed window (Figure 9-35). Select **Clone This System for Autoyast** (see “SLES 11 automated installation” on page 395 for more information) if desired and then select **[Finish]**.

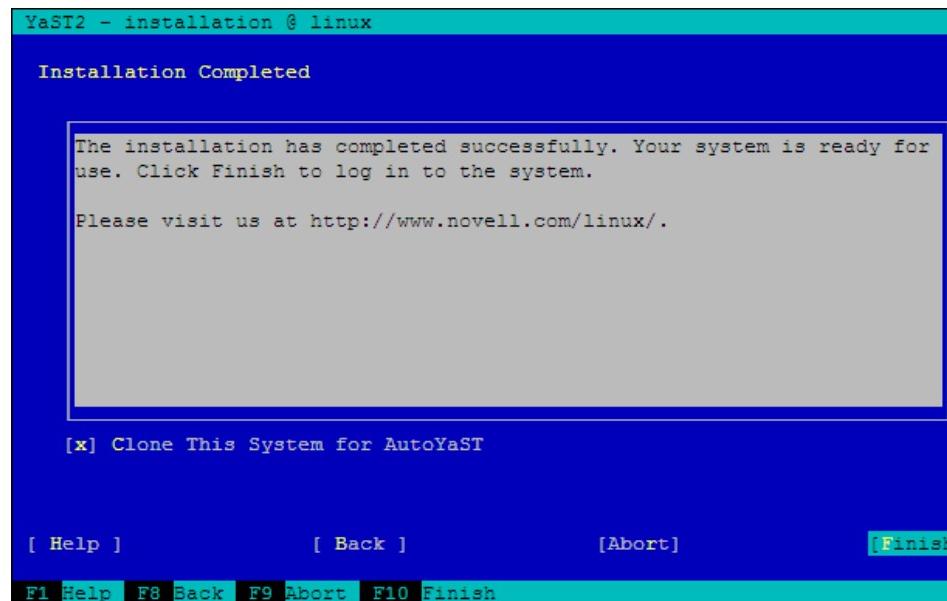


Figure 9-35 Installation completed window

29.Login to the system with the new user., as shown in Figure 9-36 on page 394.

```
Method:          3 (Parameter/Value Pairs via stdin)
Match:          serviceable=1 and (type=2 or type=3)

Servicelog ID:    4
Log Timestamp:   Sun Mar  8 14:21:21 2009
Update Timestamp: Sun Mar  8 13:21:21 2009
Notify:          0 (EVENT)
Command:         /etc/ppc64-diag/ppc64_diag_migrate
Method:          3 (Parameter/Value Pairs via stdin)
Match:          refcode="#MIGRATE" and serviceable=0

Starting rtas_errd (platform error handling) daemon:          done
Starting Firewall Initialization (phase 2 of 2):          done
Master Resource Control: runlevel 3 has been               ready
Skipped services in runlevel 3:          smbfs nfs smartd splash

Welcome to SUSE Linux Enterprise Server 11 (ppc64) - Kernel 2.6.27.19-1
onsole).

sles11 login: [REDACTED]
```

Figure 9-36 Login screen

9.5 Native SLES 11 installation

A native SLES 11 installation of a JS43 blade follows a similar process to those given in the VIOS LPAR installation sections. However, there are some key differences:

- ▶ In a native installation, the IVM terminal is no longer available to complete the Linux installation, but you can use the Serial Over LAN (SOL) console as an alternative. See Appendix A, “Consoles, SMS, and Open Firmware” on page 493 for more information. Use the SOL console to display the SMS menu and the yaboot/lilo/grub or yast options during the installation.
- ▶ The resource allocation of processors, I/O adapters, memory, and storage devices in a native environment is *fixed*.
- ▶ Virtualization functions and features are not available.

9.6 SLES 11 automated installation

SuSE has an automated installation functionality known as Autoyast to install multiple systems in parallel. The system administrator performs an Autoyast automated installation by creating a single file containing answers to all the questions normally asked during a SuSE installation. This file resides on a single server system and multiple clients can read it during installation. There are multiple methods to create an Autoyast profile:

- ▶ Clone the install configuration information from a reference machine.
- ▶ Use the Autoyast GUI to create and modify the Autoyast profile. See Appendix B, “SUSE Linux Enterprise Server AutoYaST” on page 521 for more information on this method.
- ▶ Use an XML editor to create an Autoyast profile from scratch.

This section describes a method to clone installations to identical machines from a reference machine:

1. Perform a CD/DVD or network installation.
2. In the Installation Completed YaST window (see Figure 9-35 on page 393), select **Clone This System for Autoyast** and then select **[Finish]**. This creates an Autoyast profile in /root/autoinst.xml that is ready for immediate use.

Tip: *SUSE Linux Enterprise Server 11 Doc* is a good source of information regarding all Autoyast profile options. It is available at:

http://www.novell.com/documentation/sles11/pdfdoc/book_sle_deployment/book_sle_deployment.pdf Part IV. Automated Installations

3. Place the autoinst.xml file on a server that is accessible for all the systems with *identical* hardware configurations to use.

Performing an Autoyast installation

Pass the location of the Autoyast profile and install images during boot. BOOTP does not have the ability to provide anything more than the location to the bootable image and the server IP address hosting the Autoyast profile. Use the Open Firmware prompt to pass the required parameters.

1. Type 8 during the LPAR boot process to go to the Open Firmware prompt, as shown in Figure 9-37.

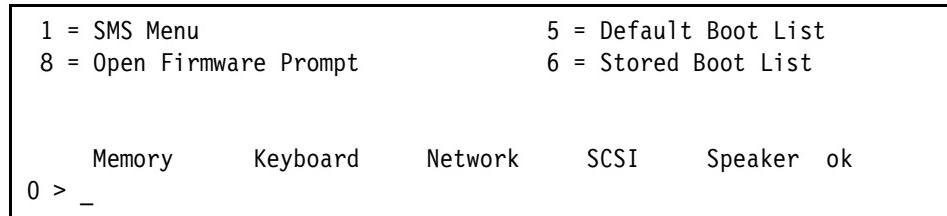


Figure 9-37 Open Firmware prompt

2. Type the following command in the Open Firmware prompt to start automated installation. For example, if the profile is served using NFS:

```
boot net autoyast=nfs://193.200.1.80/home/autoinst.xml  
install=nfs://192.168.1.254/SuSE/SLES11
```

Note: This automated install assumes there is a DHCP server that will provide the client with the correct IP address information. A system using static IP will require additional parameters for sshpassword, hostip, netmask, gateway, and nameserver.

3. Press the Enter key to start the process. The automated SUSE Linux Enterprise Linux installation is now complete.

9.7 IBM service and productivity tools

The IBM service and productivity tools are packages that are installed after a successful SLES installation.

Important: These packages are *not* shipped with the SUSE installation CDs. They are IBM owned and distributed.

These packages enable features such as:

- Reliability, availability, and serviceability (RAS) functionality
- I/O hotplug
- Dynamic Logical Partitioning (DLPAR) capabilities
- Live partition migration capabilities

See Appendix D, “Service and productivity tools for Linux” on page 545 for more information on installing the service and productivity tools specific to your system’s configuration.



JS23 and JS43 power management using EnergyScale technology

The EnergyScale technology described in 3.4, “IBM EnergyScale technology” on page 47 can be used by the BladeCenter Advanced Management Module and Active Energy Manager (AEM) to monitor and control power usage of the IBM BladeCenter JS23 and JS43 blades. This chapter describes how to use the BladeCenter AMM and Active Energy Manager extension of IBM Systems Director to utilize these features.

An in-depth discussion of installing, configuring, and using Active Energy Manager can be found in *Going Green with IBM Active Energy Manager*, REDP-4361.

EnergyScale options can be performed on the AMM and the BladeServers JS23/JS43 using AMM options or Active Energy Manager.

This chapter has the following sections:

- ▶ “Power management through the AMM” on page 398
- ▶ “Power management through AEM” on page 407

10.1 Power management through the AMM

The IBM BladeCenter Advanced Management Module (AMM) provides a Web-based and command line user interface to monitor and control individual blades and switch modules installed in the BladeCenter. The AMM also collects historical or trend data for individual components in the IBM BladeCenter. This data can be reviewed from the user interface. The information can also be collected by the Active Energy Manager extension for IBM Systems Director. This section describes how to use both interfaces of the AMM to make changes in the Power Saving and Power Capping modes. The examples shown here are from a BCH. The BCS screens look the same with the exception of having just one power domain in the BCS.

10.1.1 Using the AMM Web UI for blade power management

A detailed description of how to access the AMM will not be covered in this section but detailed access and use information can be found in *IBM eServer BladeCenter Systems Management*, REDP-3582. The Web UI can be used from a Web browser pointed to the IP address of the AMM.

AMM Power Management View

From the main menu panel in the **Monitors** section, click **Power Management**; see Figure 10-1 on page 399 for an example of the Power Management option. The BladeCenter H (BCH) has two power domains. The BladeCenter S (BCS) only has one power domain.

The AMM Power Management view is divided into the following sections:

- ▶ BladeCenter Power Domain Summary
- ▶ BladeCenter Power Domain Planning
- ▶ BladeCenter Chassis Power Summary
- ▶ BladeCenter Chassis Configuration Setting
- ▶ BladeCenter Chassis Power Consumption

In the graphic below there are 2 highlighted boxes. Number 1 links to the AMM Power Management Policy. Number 2 links to the Power Domain which contains BladeServer slots and chassis components.

In each section a complete description of the fields can be obtained by clicking the blue question mark.

BladeCenter Power Domain Summary

		Power Domain 1	Power Domain 2
Status	Power domain status is good.		
Power Modules	Bay 1: 2880W Bay 2: 2880W	Power domain status is good. Bay 3: 2880W Bay 4: 2880W	
Power Management Policy	1. Power Module Redundancy with Blade Throttling Allowed Very similar to Power Module Redundancy. This policy allows you to draw more total power; however, capable blades may be allowed to throttle down if one Power Module fails.		
Maximum Power Limit [*]	3380W	3380W	
Power in Use ^{**}	200W	1298W	

BladeCenter Power Domain Planning

		Power Domain 1	Power Domain 2
Maximum Power Limit [*]	3380W	3380W	
- Allocated Power (Max) ^{***}	1195W	2280W	
= Remaining Power	2185W	1100W	

* Maximum power available based on the number of power modules and the Power Management Policy setting.
** Represents the maximum worse case and measured power based on the capability of all components.
*** Reserved power for all components in this domain.

BladeCenter Chassis Power Summary

Total DC Power Available	6760W
Total AC Power In Use [*]	2018W
Total Thermal Output	6,885.4 BTU/Hour

* Includes Chassis Cooling Devices that are AC powered.

Figure 10-1 BladeCenter Power Domain Summary

Scrolling the page down below the Blade Chassis Power Summary will provide access to the acoustical settings for the chassis, power consumption history and links to view the thermal and power trending history for some of the chassis components. An example of the options is shown in Figure 10-2 on page 400 and Figure 10-3 on page 400.

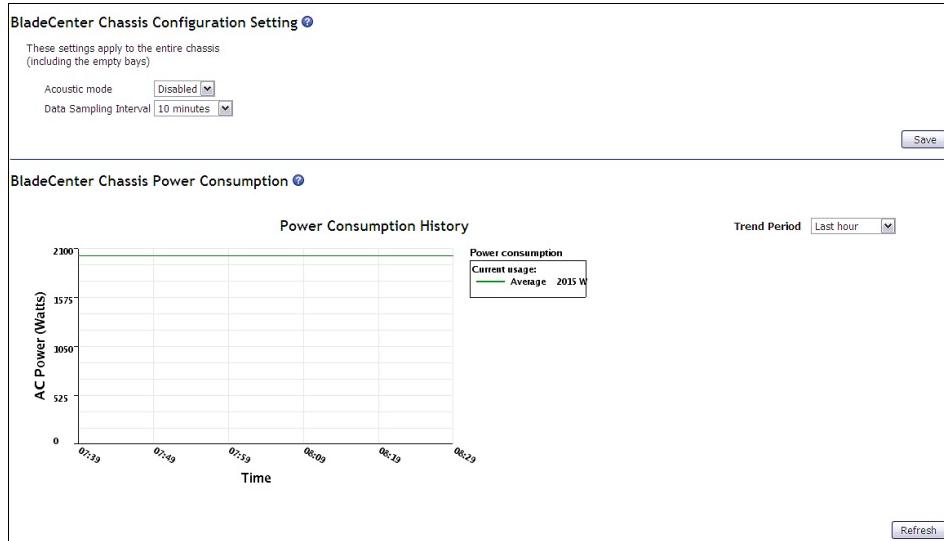


Figure 10-2 Additional power settings

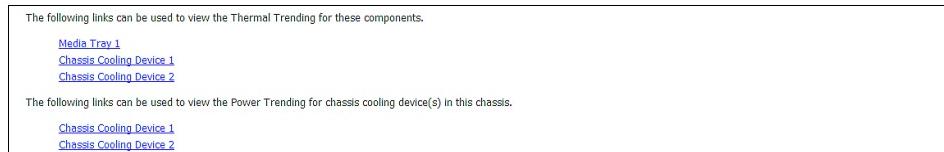


Figure 10-3 Chassis thermal and trending options

Selecting the **Power Management Policy** link (number 1 as shown above) will allow the user to select three different management policies. Figure 10-4 on page 401 shows an example of this option.

There are three different selections that can be applied to manage the power domain. As mentioned above, in the BCH there are two power domains. Each domain can set this policy separately and they do not need to match.

The Basic Power Management option is the least conservative of the three available options. This policy does not guarantee redundancy so if a power module fails, the BladeCenter chassis could be affected. When this policy is in effect the value for Maximum Total Power Limit is used to determine if a blade can power on. As long as the power used is below this maximum value the blade can continue the power on procedure. If one of the power modules fail, blades capable of throttling may be throttled back to keep the domain usage below the maximum limit. If the limit is reached, the domain may be affected by power loss.

The Power Module Redundancy option is used when only one AC source is present. One AC source in this case means the electrical grid. For example, the BCH has two line cord inputs. Each is capable of connecting to its own AC power source. If the two line cords attach to the same power grid, it is considered a single AC source. It is possible to have a data center wired so that each AC line cord of the BCH could be plugged into a separate power grid or AC source. If an outage occurs on one power module, this option allows some redundancy and may keep the chassis and blades running. A second power module failure may result in a domain outage.

The third option, Power Module Redundancy with Blade Throttling is basically the same as the Power Module Redundancy option except in case of a failure of a module, the blades can be throttled back to allow the chassis and blades to continue to run. Once an outage has been restored, the blades will return to running at the un-throttled power level.

BladeCenter Domain 1 Power Management Policies 				
Links	Power Summary	Domain 2 Power Management Policies		
Select	Option Name	Power Supply Failure Limit [†]	Maximum Power Limit (Watts)	Estimated Utilization ^{††}
<input type="radio"/>	Power Module Redundancy Intended for a single AC power source into the chassis where each Power Module is on its own dedicated circuit. Total allowed power draw is limited to one less than the number of Power Modules when more than one Power Module is present. One Power Module can fail without affecting blade operation. Multiple Power Module failures can cause the chassis to power off. Note that some blades may not be allowed to power on if doing so would exceed the policy power limit. More...	1	2880	6%
<input checked="" type="radio"/>	Power Module Redundancy with Blade Throttling Allowed Very similar to Power Module Redundancy. This policy allows you to draw more total power; however, capable blades may be allowed to throttle down if one Power Module fails. More...	1	3380	5%
<input type="radio"/>	Basic Power Management Total allowed power is higher than other policies and is limited only by the total power capacity of all the Power Modules up to the maximum of chassis power rating. This is the least conservative approach, since it does not provide any protection for AC power source or Power Module failure. If any single power supply fails, blade and/or chassis operation may be affected. More...	0	3520	5%

[†] This is the maximum number of power supplies that can fail while still guaranteeing the operation of the domain in the selected policy. ^{††} The estimated utilization is based on the maximum power limit allowed in this policy and the current aggregated power in use of all components in the domain.

[Save](#) [Refresh](#)

Figure 10-4 Power Management Policies

Selecting the Power Domain link as referenced by number 2 in Figure 10-1 on page 399 will allow the user to select different components of the chassis to modify the power settings. This is where you can select the BladeServer JS23/JS43 to modify the power policies.

AMM Power domain details view

The BladeCenter Power Domain Details view shows each module that is serviced by the power domain. The chassis location (bay), status, module type, power state, current power in use, allocated maximums and minimums in watts, and CPU duty cycles are shown for each module. Modules that have specific

capabilities or collect power trend data appear as a link to a module-specific detail view. Figure 10-5 on page 402 provides an example of this selection.

BladeCenter Power Domain 2 Details 						
Bay(s)	Status	Module	State	Power In Use	Allocated Power	CPU Duty Cycles
<i>Chassis Components</i>						
	Midplane		On	5W	10W	10W
1	Media Module		On	5W	10W	n/a
<i>Power Module Cooling Devices</i>						
1	Power Module		On	15W	30W	n/a
2	Power Module		On	15W	30W	n/a
3	Power Module		On	15W	30W	n/a
4	Power Module		On	15W	30W	n/a
<i>Management Modules</i>						
1	xbch4amm		On	12W	25W	n/a
2	Advanced Management Module Bay 2 (not present)			0W	15W	n/a
<i>I/O Modules</i>						
1	Ethernet SM		On	22W	45W	n/a
2	Ethernet SM		On	22W	45W	n/a
7	Fibre Channel SM		On	15W	30W	n/a
9	Fibre Channel SM		On	32W	65W	n/a
<i>Blades</i>						
[8]	JS23-Redbook		On	213W	406W	n/a **
[9-10]	JS43-Redbook		On	462W	745W	n/a **
[13-14]	Japan_demo		On	456W	764W	n/a **

 * This blade may throttle if redundancy is lost in this power domain.
 ** Click on the module name to view CPU speeds.
 * Cannot communicate with the blade. The power values for this blade are assumed.

Figure 10-5 Power Domain Details

Selecting the components such as a blade will allow you to set some of the power management options. Shown below in Figure 10-6 on page 403 you can see the options available for a blade that is capable of power management.

In this panel you can see what the blade power capabilities are. In this example the blade supports:

- Power metering
- Power capping
- Static low power saving
- Dynamic power saving

The screenshot shows two related web pages for managing power settings on a BladeCenter JS23-Redbook system.

BladeCenter JS23-Redbook (Bay 8) Power Summary

- The following capabilities are supported:**
 - Power metering
 - Power capping
 - Static low power saving
 - Dynamic power saving

Processors	4
Effective CPU Speed	3621 MHz
Maximum CPU Speed	4204 MHz

BladeCenter JS23-Redbook (Bay 8) Configuration Setting

Power Capping Options

- Power Capping:** Enabled (dropdown menu)
- Maximum Power Limit (range 180-406, guaranteed range 320-406):** 406

Power Savings:

- Static High Performance
- Static Low Power Saver
- Dynamic Power Saver
- Favor Performance over Power

Save

Figure 10-6 Blade power configuration settings

Power capping is used to allow the user to allocate less power and cooling to a system. This can help save on datacenter infrastructure costs, and then potentially allow more servers to be put into an existing infrastructure.

To enable the *Power Capping* option, use the pull down menu and select *Enable*. Then you will be able to set a cap level using the *Maximum Power Limit* range box. This value will limit the power usage to the value specified. When the limit has been reached the blade will be throttled back to limit the consumption to the maximum value.

Power savings is used to put the server into a mode that consumes less energy.

Power savings can be enabled using one of the various options. Selecting the *Dynamic Power Saver* also permits you to select the option to favor performance over power usage.

Scrolling the page down will allow the user to view trend data for the particular bladeserver selected. Figure 10-7 on page 404 shows an example of the trend data.

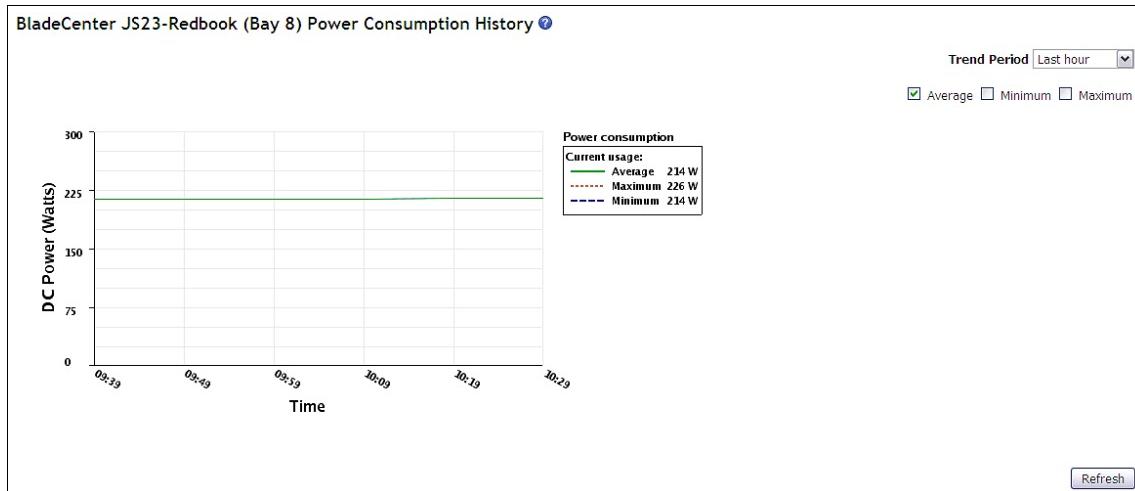


Figure 10-7 Bladeserver trend data

10.1.2 Using the AMM CLI UI for blade power management

Similar to the Web UI, the CLI can be used to display power domain and specific module information. The AMM CLI can be accessed by either a telnet or SSH to the IP address of the AMM. The login is completed by using the same user ID and password that is used for the Web UI.

The fuelg command

The **fuelg** command used from the AMM CLI controls and displays BladeCenter power functions and the EnergyScale features of an IBM BladeCenter JS23 or 43. This command with all options is shown in Example 10-1.

Example 10-1 fuelg command with help flag

```
system> fuelg -h
fuelg {<pdx> {-pm} |{-pt}}|{-am|-int|-pcap|-pme|-ps}|{-pt}|{-tt}
Power management settings and values by domain
pdx:   power domain (pd1, pd2)
-pm:   power management settings (redwoperf, redwperf, nonred)
-am:   acoustic mode setting (on, off)
-int:  polling interval (between 10 and 60 minutes, in increments of
5)
-pcap: power cap for blades. Note: setting a pcap will automatically
enable
pme.
```

-pme: power management and capping enabling for blades (off, on).
Note: the
blade must be powered on before enabling capping.
-ps: power saver mode for blades (off, on). Note: the blade must be
powered
on before enabling power saver mode.
-pt: power trending data (1, 6, 12, or 24 hours)
-tt: thermal trending data (1, 6, 12, or 24 hours)

Example 10-2 shows the **fuelg** command used from the system> prompt with no flags to display the BladeCenter Power Domain information.

Example 10-2 The fuelg command used to display BladeCenter Power Domain

system> fuelg
Note: All power values are displayed in Watts.

Total DC Power Available: 2900
Total AC Power In Use: 807
Total Thermal Output: 2753 BTU/hour

Power Domain

Status: Power domain status is good.
Modules:
Bay 1: 1450
Bay 2: 1450
Bay 3: 1450
Bay 4: 1450
Power Management Policy: AC Power Source Redundancy
Power in Use: 769
Total Power: 2900
Allocated Power (Max): 1346
Remaining Power: 1554
Power Service: 220 VAC

-am off
-int 10

From the AMM system> prompt use the **env** command to set the specified blade as the persistent target for commands during the current session. Example 10-3 shows the environment being set to blade slot 4. Note the prompt change.

Example 10-3 The env command used to set a persistent target

```
system> env -T blade[4]
OK
system:blade[4]>
```

Example 10-4 shows the **fuelg** command with no other parameters being used to display the capabilities, current settings, and power consumption values of the blade in BladeCenter slot 4.

Example 10-4 The fuelg command used to display blade slot power information

```
system:blade[4]> fuelg
-pme off
-ps off
PM Capability: Dynamic Power Measurement with capping and power saver mode
Effective CPU Speed: 3800 MHz
Maximum CPU Speed: 3800 MHz
-pcap 256 (min: 256, max: 282)
Maximum Power: 150
Minimum Power: 150
Average Power: 150
```

The **fuelg** command flags for changing blade specific EnergyScale parameters are:

- ▶ **-int *interval*** - polling interval between 10 and 60 minutes in increments of 5.
- ▶ **-pcap *value*** - power cap for blades, the value between the min and max amounts, setting pcap also enables pme.
- ▶ **-pme *on/off*** - power management and capping enabling/disabled uses min value unless a different pcap value has been used.
- ▶ **-pt *interval*** - power trending data 1,6,12, or 24 hours.

Power Saver Mode enabled from the AMM CLI

The baseline values and power consumption are shown in Example 10-4 on page 406. Power Saver Mode was enabled using the **fuelg -ps** command and then the changes reviewed using the **fuelg** command shown in Example 10-5.

Example 10-5 Power Saver Mode enabled

```
system:blade[4]> fuelg -ps on
OK
system:blade[4]> fuelg
-pme off
```

```
-ps on
PM Capability: Dynamic Power Measurement with capping and power saver mode
Effective CPU Speed: 3440 MHz
Maximum CPU Speed: 3800 MHz
-pcap 256 (min: 256, max: 282)
Maximum Power: 139
Minimum Power: 139
Average Power: 139
```

Power trend data for the last hour was reviewed using the **fuelg -pt 1** command shown in Example 10-6.

Example 10-6 Power trend data in Power Saver Mode

```
system:blade[4]> fuelg -pt 1
Date      Time      Max  Min  Avg
-----
07/02/08 13:20:20  162  150  150
07/02/08 13:30:20  154  150  150
07/02/08 13:40:20  154  138  147
07/02/08 13:50:19  150  138  139
07/02/08 14:00:20  142  138  139
07/02/08 14:10:20  142  138  139
```

10.2 Power management through AEM

The Active Energy Manager (AEM) Plug-in for IBM Systems Director provides the ability to monitor and manage energy capabilities of a resource. In addition, AEM provides management functions through integration with IBM Systems Director such as scheduling options to apply power saver mode and power capping options to support systems or groups of systems at predetermined times.

AEM also supports the application of power policies to supported systems or groups of systems thereby managing energy across multiple systems can be deployed, maintained, and modified with minimal effort. IBM Systems Director Version 6.1.1 and Active Energy Manager version 4.1.1 are required for energy management of the JS23/JS43 BladeServers

It is not the intent of this redbook to explain and demonstrate all the different options available using AEM. The references in this publication only intend to show some of the options available and what can be configured. There are in

most instances multiple paths to the same options in AEM. The AEM redbook that will be created shortly after this publication will provide details on these options in greater detail.

The following information and examples assume that IBM Systems Director and the Active Energy Manager extension have been installed and configured. Complete planning, installation, configuring, and usage information of IBM Systems Director can be found in:
www.redbooks.ibm.com/redpieces/abstracts/sg247694.html.

IBM Active Energy Manager installation and use can be found in *Going Green with IBM Active Energy Manager*, REDP-4361. At the time of this writing the above guides may not have been updated to the new version. It is assumed they will be updated at some point after this redbook has been published. The links should guide you to the newest information.

10.2.1 Active Energy Manager Console

Once IBM Systems Director has been installed and your target systems have been configured you can use AEM to manage the power capabilities of the chassis and bladeservers.

To access the Active Energy Manager option, use the menu option as shown in Figure 10-8 on page 409.

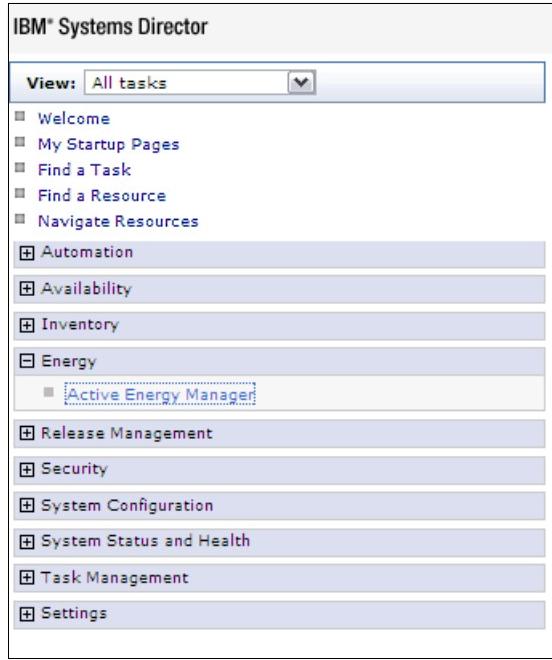


Figure 10-8 Director menu options

Once AEM has been selected you will have the options available in Figure 10-9 on page 410. In this example we have four resources that can be managed by AEM. One of the resources is the BCH chassis and the other three are bladeservers within the chassis.

Note: When a JS43 is present in the chassis, the AMM may have problems reporting the JS43 BladeServer to AEM. To correct this issue be sure that the AMM firmware level is at BPET48F or higher. Otherwise the JS43 may not appear as a resource that can be managed by AEM.

The screenshot shows the Active Energy Manager interface. The top navigation bar includes 'Active Energy Manager' and a 'Settings' link. Below the header, a message states: 'Work with power-managed resources. View recent power and temperature status. Monitor power and environmental values. Configure power settings and automate tasks in response to power and environmental events.'

Status

Top 5 highest average input power values

	Last 30 days
Today	1,937W xbch4amm 464W IBM 7778 63X 10194FA 412W IBM 7778 63X 10181CA 232W IBM 7778 62X 10180EA 147W Cooling Module 1

Top 5 highest ambient temperature values

	Last 30 days
Today	73.4F 71.6F xbch4amm

Monitor

Navigate the list of resources managed by Active Energy Manager. Right-click on a resource to view properties and perform tasks.

To launch a full-sized view of the table in a new tab, choose the link directly above the table.

Active Energy Manager Resources (View Members)

Actions	Name	Type	Description	Average Input Power	Average Output	Ambient Temp	Exhaust Temp
<input checked="" type="checkbox"/>	IBM 7778 62X 10180EA	Server	JS23 BladeServer	213.041			
<input type="checkbox"/>	IBM 7778 63X 10181CA	Server	JS43 BladeServer	397.547			
<input type="checkbox"/>	IBM 7778 63X 10194FA	Server	JS23 BladeServer	455.979			
<input type="checkbox"/>	IBM 7998 61X 100180A	Server	JS22 BladeServer	8.607			

Page 1 of 2 | Selected: 1 Total: 6 Filtered: 6

Figure 10-9 Active Energy Manager options

10.2.2 AEM Energy Properties

Using the check box you can select the resource to work with. Figure 10-10 on page 411 shows an example of selecting the BladeCenter Chassis and then using the Actions button to select the **Properties** option as displayed in Figure 10-11 on page 411.

Active Energy Manager Resources (View Members)				
Actions	Search the table...		Search	
Select	Name	Type	Description	
<input type="checkbox"/>	IBM 7778 62X 10180EA	Server		
<input type="checkbox"/>	IBM 7998 61X 100180A	Server		
<input type="checkbox"/>	IBM 7998 61X 1001C8A	Server		
<input checked="" type="checkbox"/>	xbch4amm	BladeCenter Chass		

Figure 10-10 Select resource



Figure 10-11 Actions options

Using the various tabs properties view you can see information about the resource selected. Clicking on the Active Energy tab allows you to view the data available about the chassis as shown in figure Figure 10-12 on page 412.



Figure 10-12 Properties - Active Energy tab

Using the Edit tab you can modify the energy price and metering values. This data can then be used for cost estimating of the power used for the chassis. Figure 10-13 on page 412 shows an example of the values available to edit.

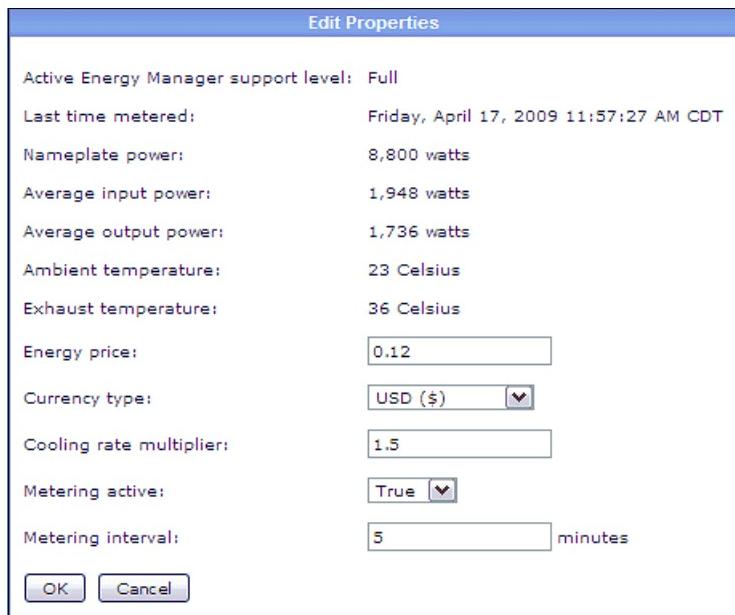


Figure 10-13 Edit values

10.2.3 BladeCenter Energy Properties

In this next section we will look at the energy management options available on the JS23/JS43. Using AEM you can configure power capping, power saver and view trend data on the bladeserver.

Enabling Power Capping

To enable power capping on the bladeserver use AEM and select the desired blade resource. Using the Actions button select **Energy** then **Manage Power** and finally **Power Capping** as shown in Figure 10-14 on page 413.

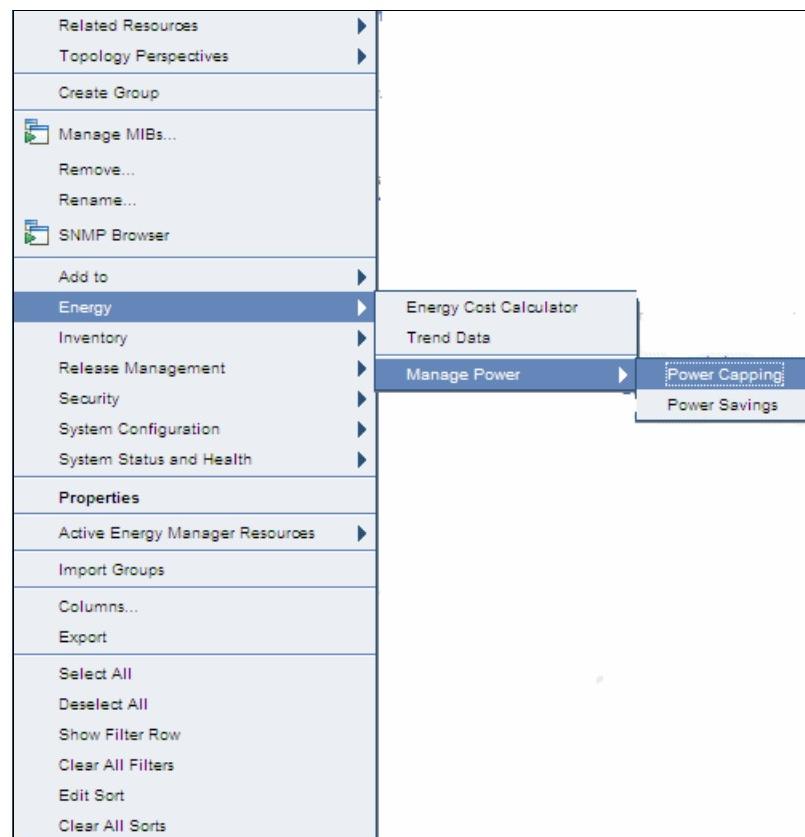


Figure 10-14 Power Capping option

To activate Power capping, select the **Activate Power Capping** option. Once selected you can modify the parameters for power cap type and set the value. There are two options for the power cap type. One is based on a percentage, the other option is based on wattage. Choose the preferred value and set the value using the slider bar or input a number in the entry box. Click on the **Save** button

to save your settings. An example of the power capping options are shown in Figure 10-15 on page 414.

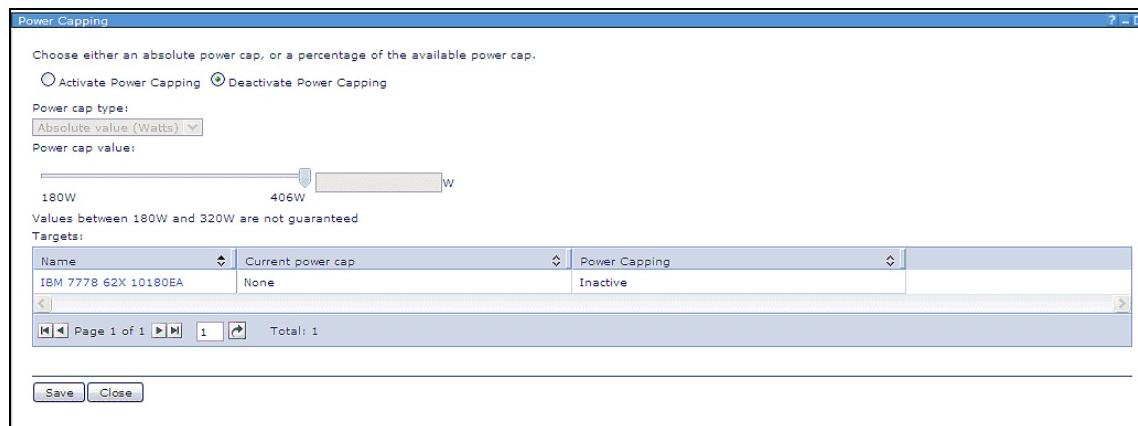


Figure 10-15 Power Capping options

Figure 10-16 on page 414 shows an example of the power capping features enabled for the bladeserver.

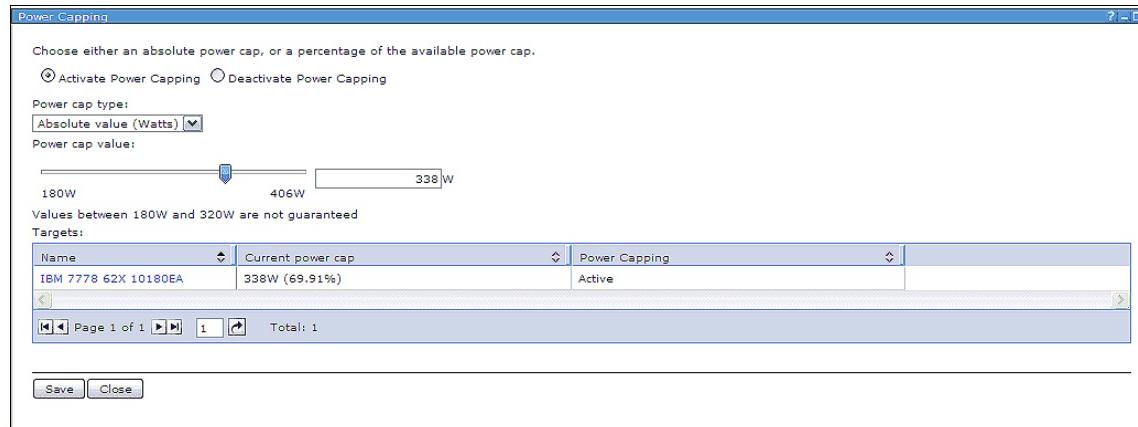


Figure 10-16 power capping enabled

Enabling Power Savings

To enable power savings on the bladeserver use AEM and select the desired blade resource. Using the Actions button select **Energy** then **Manage Power** and finally **Power Savings** as shown in Figure 10-17 on page 415.

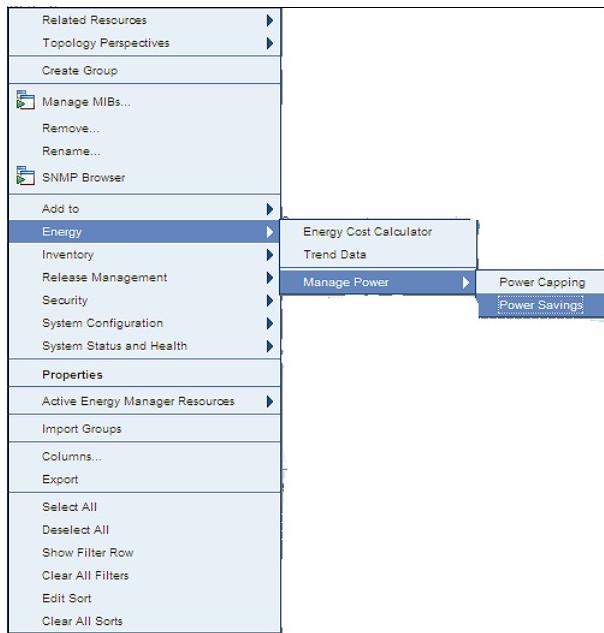


Figure 10-17 Power Savings option

The power savings options are as follows:

No power savings - choose this option to have no power savings. The processor runs at high speed.

Static power savings - choose this option to reduce power usage by lowering processor speed. This option saves energy while maintaining a reasonable processor performance.

Dynamic power savings - choose this option to automatically balance power usage and processor performance. This option saves energy while minimizing performance impact. When dynamic power savings is enabled, you can also select to **Favor power** or **Favor performance**.

Figure 10-18 on page 416 shows an example of the power savings options. Click on the **Save** button to save your modified preferences.

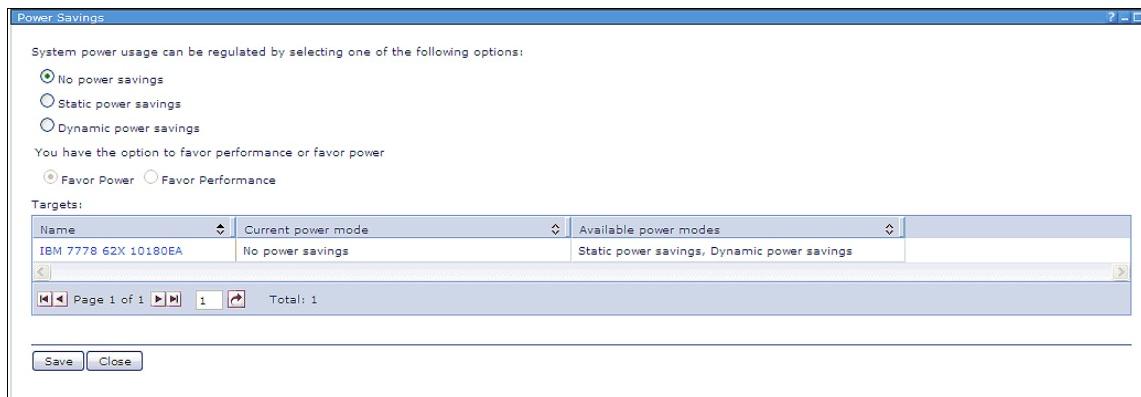


Figure 10-18 Power Savings options

Viewing BladeServer JS23/JS43 Trend Data

Using the AEM you can view trend data for the JS23/JS43. Trend data provides information usable to view details relating to power usage, capping values and informational events. This data can be charted for the last hour up to the last year in different intervals. Figure 10-19 on page 416 shows an example of selecting the **Trend Data** details.

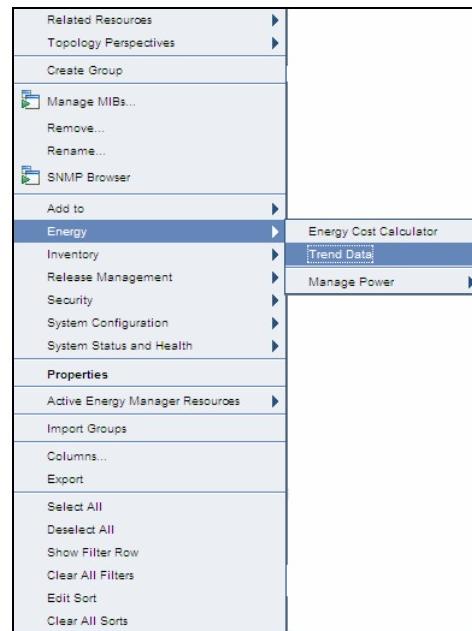


Figure 10-19 Trend Data option



Figure 10-20 Trend Data display

In the trend data panel you can view various power details. Use the pull down menu to change the time period or click on the Custom Settings link to change the values. Click on Refresh Trend Data to see your changes.

Scrolling down in the trend data display will show information on environment data such as temperature.

Chart data can be modified as well using the Options link. This option allows you to tailor the details on the trend data output based on the data you wish to track. Figure 10-21 on page 418 shows an example of the options available.

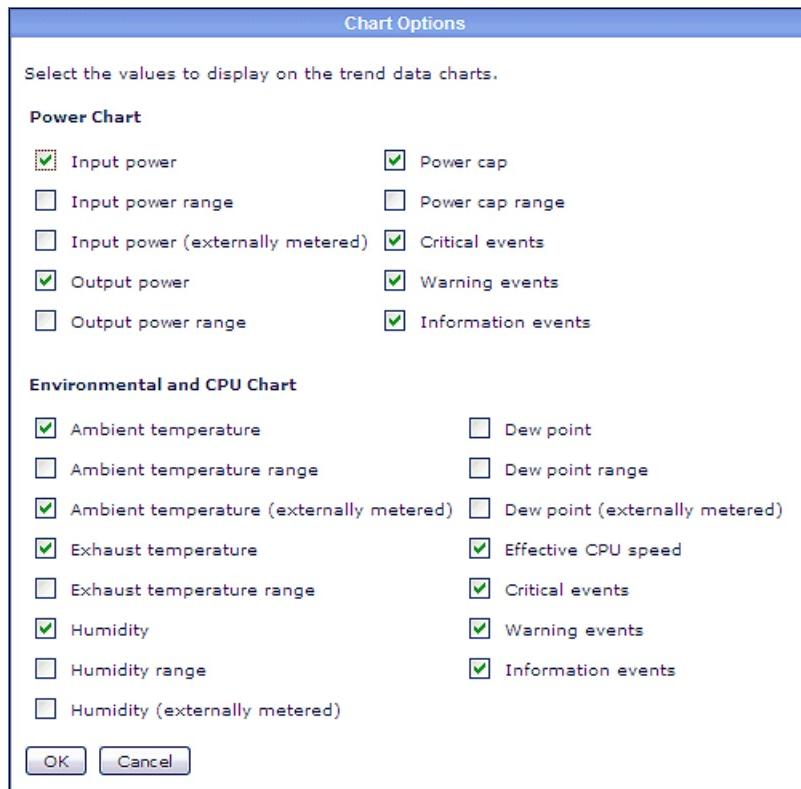


Figure 10-21 Trend data chart options

Information events as noted by the icon will display details about the event if you mouse over the icon. In this example in Figure 10-22 on page 418 you can see that a mode change was made on a resource.

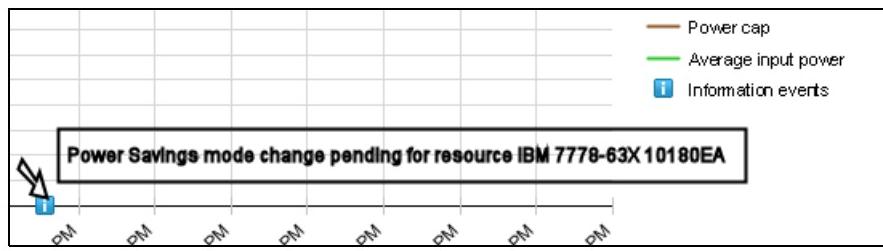


Figure 10-22 Information event details

Trend data may also be exported to your Director Server file system. Use the export option and save the file in your preferred location. Figure 10-23 on

page 419 provides an example of this option. The file is then viewable using a spreadsheet program like Excel®.

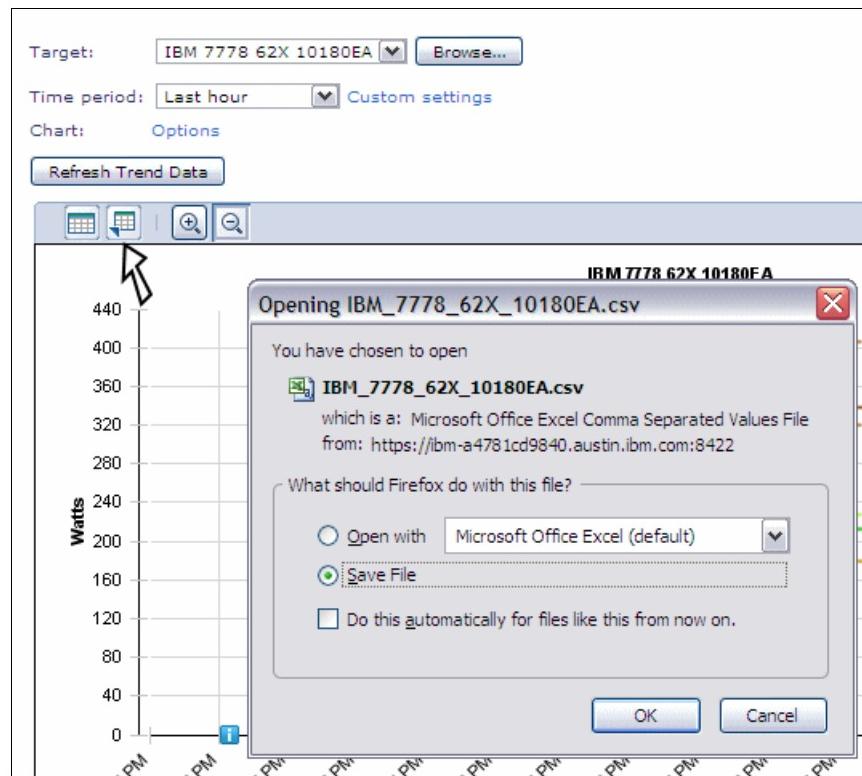


Figure 10-23 Export data

Energy Cost Calculator

Active Energy Manager has a calculator that can help determine the cost of energy for the monitored resource. Use the options **Energy** then **Energy Cost Calculator** to use this function. Figure 10-24 on page 420 shows the option to select.

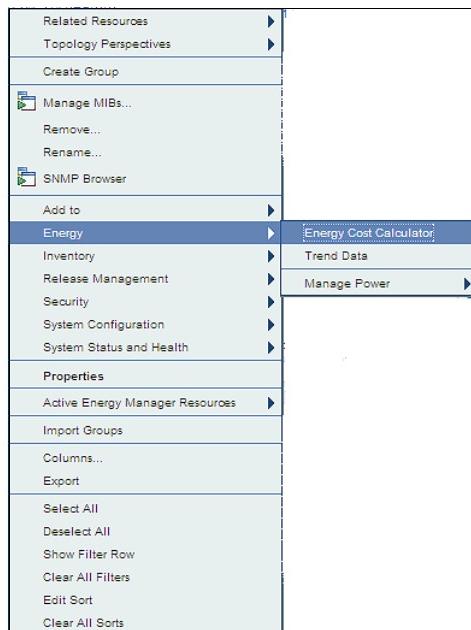


Figure 10-24 Energy calculator option

Set the values for the cost of energy using the cost properties link. Set the values for energy cost, currency type and other values. Click **OK** to save the properties. Figure 10-25 on page 420 displays an example of the properties options.

Edit Properties

Active Energy Manager support level:	Full
Last time metered:	Friday, April 17, 2009 2:46:27 PM CDT
Nameplate power:	8,800 watts
Average input power:	1,942 watts
Average output power:	1,730 watts
Ambient temperature:	23 Celsius
Exhaust temperature:	35 Celsius
Energy price:	<input type="text" value="0.12"/>
Currency type:	<input style="width: 100px; height: 20px; border: 1px solid black; border-radius: 5px; padding: 2px 10px;" type="button" value="USD (\$)"/>
Cooling rate multiplier:	<input type="text" value="1.5"/>
Metering active:	<input checked="" type="checkbox"/> True
Metering interval:	<input type="text" value="5"/> minutes

OK **Cancel**

Figure 10-25 Energy cost properties

Select the **Calculate Energy Cost** button to see the data. Figure 10-26 on page 421 shows an example of the data displayed.

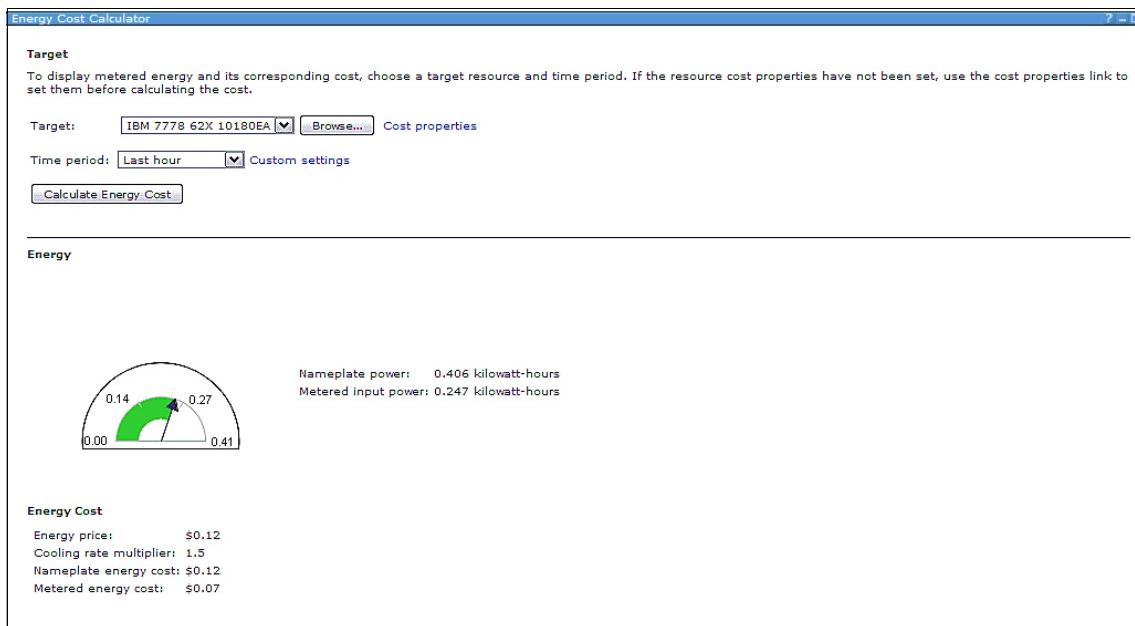


Figure 10-26 Calculated energy cost

10.2.4 Creating Power Policies

AEM supports the creation and application of power policies to manage energy across a group of systems. This feature allows you to create an energy policy and deploy the policy across a group or individual supported systems with minimal effort.

While IBM Systems Director is running, the power policies will be enforced. This is a change from previous levels of AEM. Policies can be applied to multiple systems and groups and can be applied immediately or scheduled. Within the policy management options you can view what policies are in effect and to which systems or groups the policies are applied to. Policies can also be removed using the same features used to create a power policy.

Using the main window of AEM, scroll down to the Manage section. Under the manage section is the option to **Work with power policies**. Figure 10-27 on page 422 shows an example of the selection next to the arrow.



Figure 10-27 Work with power policies

Selecting the option to **Work with power policies** brings up the screen as shown below in Figure 10-28 on page 422. From this screen you can view policies, launch a wizard to create policies, edit and delete policies. You will use this same interface to apply and remove policies once they have been created.

To begin a target or group of targets needs to be defined to configure a power policy to act on. Use the **Browse** button to begin the target selection.



Figure 10-28 Select targets - browse

Click on the check box to select the intended target or targets. After you complete your selections click on the **Add** button to add your selections. Figure 10-30 on page 423 shows an example of the targets selected and the Add box highlighted.

If you want to create a group policy, you will need to define that policy to a group using the browse option. Group policies can only be applied to a group and not individual systems. It is assumed you have created a group prior to using AEM. If a group has not been specified, you can use IBM Systems Director and create a group by navigating resources and using the Create Group wizard. Once the group has been created you can add members to the group.

An example of a group selection is provided in Figure 10-29 on page 423. To access groups, use the **Show** pull down and select **Groups**.

Otherwise, to select system resources use the **Active Energy Manager resources** selection as shown in Figure 10-30 on page 423.

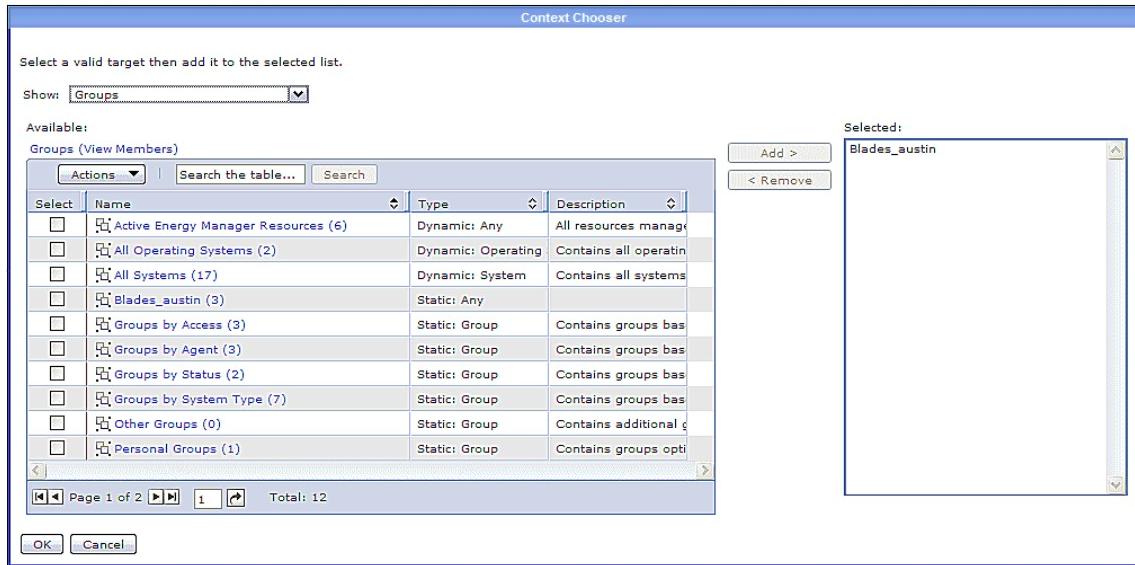


Figure 10-29 Group Select

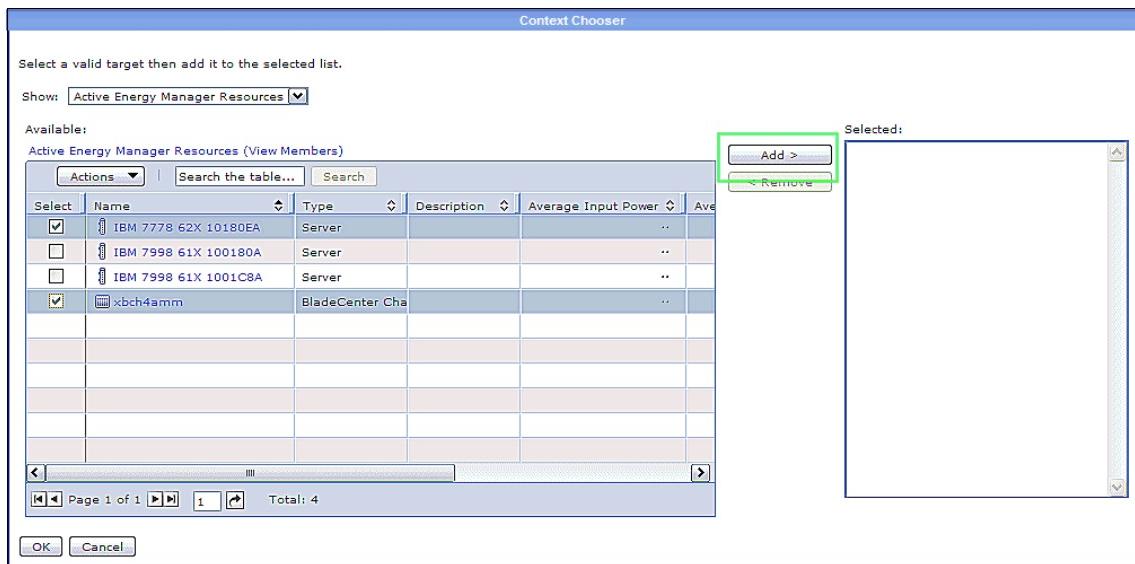


Figure 10-30 Select targets

Once your targets are added to the Selected box, click **OK** to complete your target selection. Figure 10-31 on page 424 provides an example of the targets added to the Selected box.

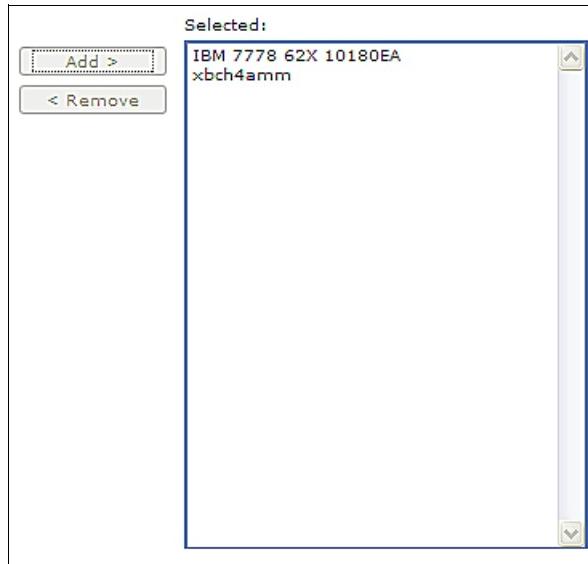


Figure 10-31 Selected targets added

Once the targets have been defined you can begin to create a power policy by clicking on the **Create Policy** button as shown in Figure 10-28 on page 422.

Clicking on the Create Policy button will start a wizard that will help you select the options for your policy. There are three different policy types that can be created. They are Group Power Capping, System Power Capping and System Power Savings. Within the policy you can select to turn on or turn off the feature.

For example: create a policy that turns on System Power Savings. Use that policy to turn on power savings for a bladeserver for use over a weekend or off-shift.

Then, create another policy that turns off power savings for the same bladeserver. Use this new policy to return the bladeserver to full power mode for week-time or on-shift usage.

Figure 10-32 on page 425 shows an example of the wizard welcome screen and a description of each policy type. To progress through the wizard, select your options, then click **Next**.

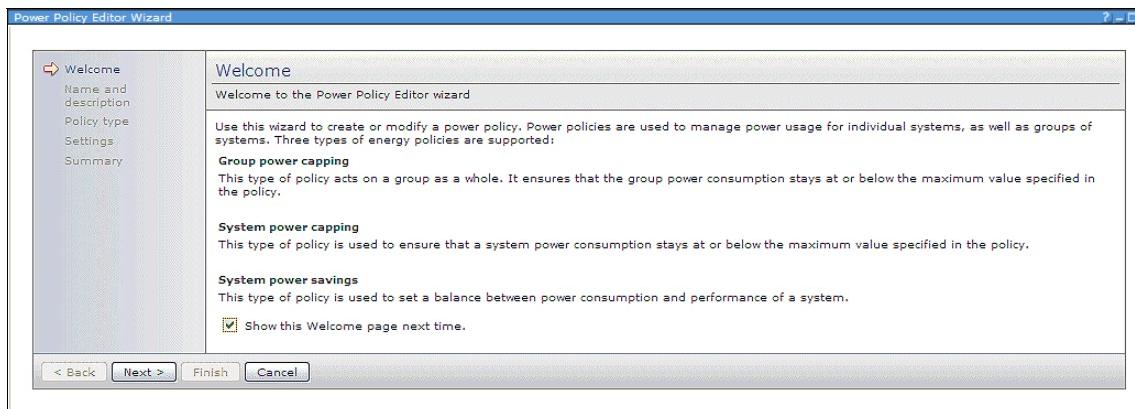


Figure 10-32 Power policy wizard welcome

In the next screen you will provide a name and description for the policy you are creating. Figure 10-33 on page 425 provides an example of this screen. The **Name** field is required, the description field is not required however, it is a good idea to describe what the policy is used for in the description field. Click **Next** to continue.

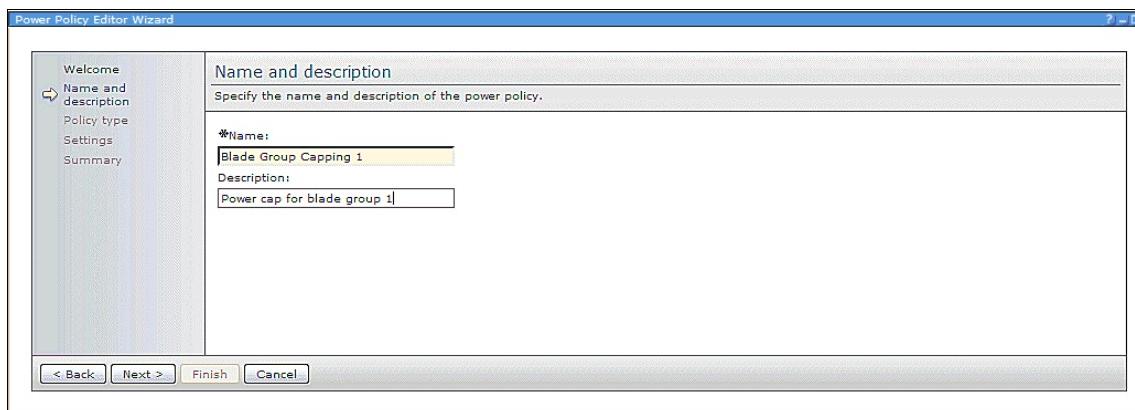


Figure 10-33 Policy name and description

In this next screen you can set the type of power policy by selecting one of the options available. The policy type screen shows the three types of policies available and also provides a short description of what the policy can do and what targets it can be used on. Figure 10-34 on page 426 shows an example of the policy type options.

For our example we are creating a power policy for a group that will provide *Group Power Capping*.

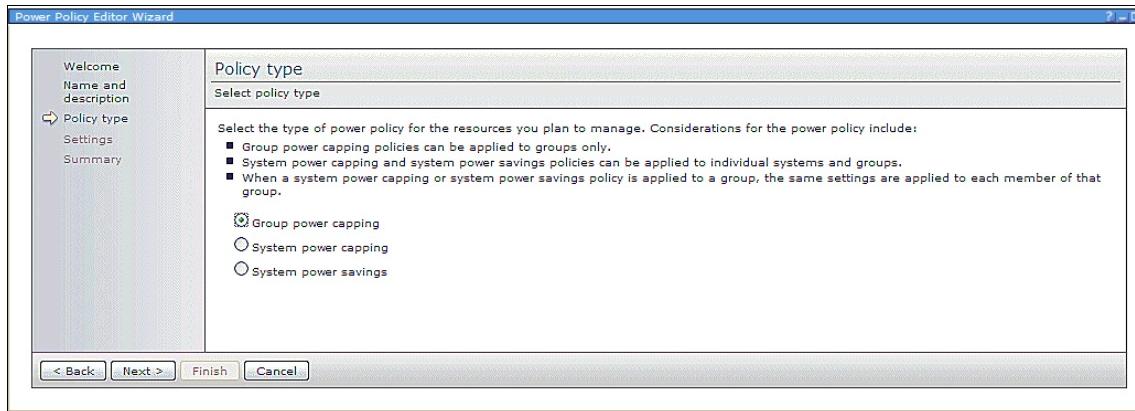


Figure 10-34 Power policy type

Select the *Group Power Capping* settings by selecting either the value in watts or use the pull down to change the value to a percentage. Set the value you wish to cap the group at in the **Power Cap Value** field. Figure 10-35 on page 426 shows an example of this screen with values for our group. Click **Next** to continue.

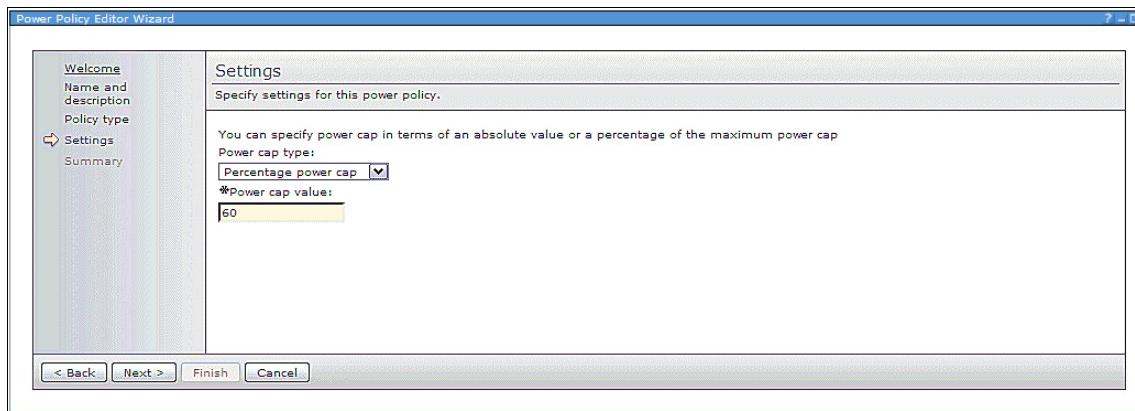


Figure 10-35 Power policy settings

The final screen of the wizard provides a summary of your selections. Verify the selections are correct and click **Finish** to complete the policy creation. Figure 10-36 on page 427 shows an example of the summary screen.

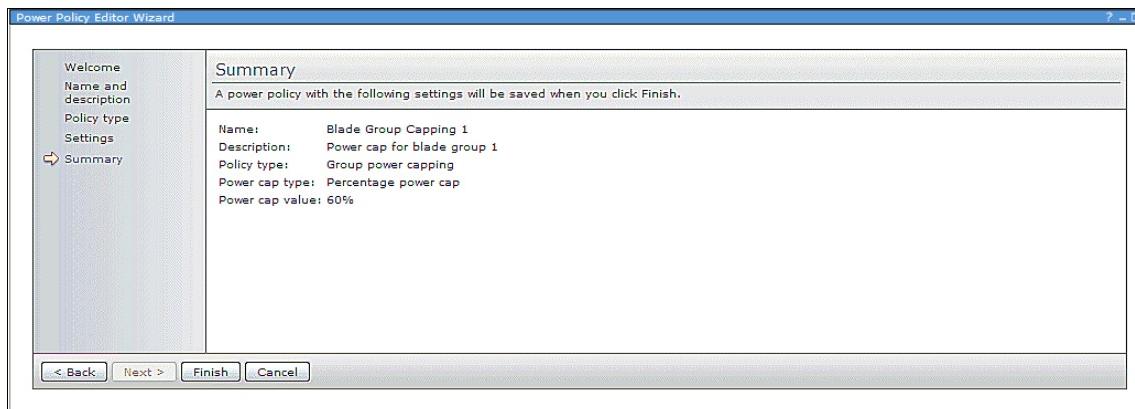


Figure 10-36 Power policy summary

Now that the policy has been created, it can be selected for action. In the next graphic you can see the policy we created with the wizard in the last few screens as well as a few other policies we created to take action on.

Figure 10-37 on page 427 provides an example of a few power policies available for actions.

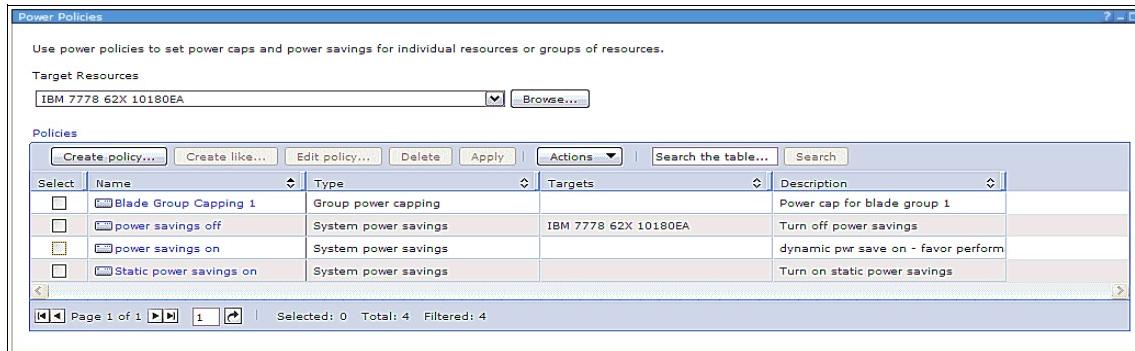


Figure 10-37 Available power policies

Now that there are power policies created we can select a target system or group of systems to apply the power policy to. Earlier we selected a group of targets using the browse and add features. To apply a power policy to our selected targets, ensure the targets are listed in the **Target resources** selection.

Then select the power policy and click on the Apply button (highlighted). Figure 10-38 on page 428 shows an example of the targets selected and the power policy to apply selected.

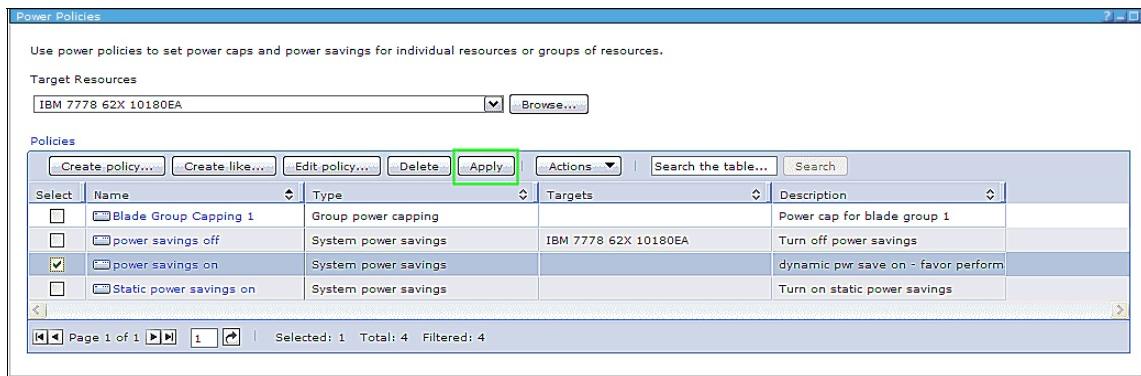


Figure 10-38 Apply power policy

In the next screen you can select when to apply the policy. Figure 10-39 on page 428 shows the apply now options.

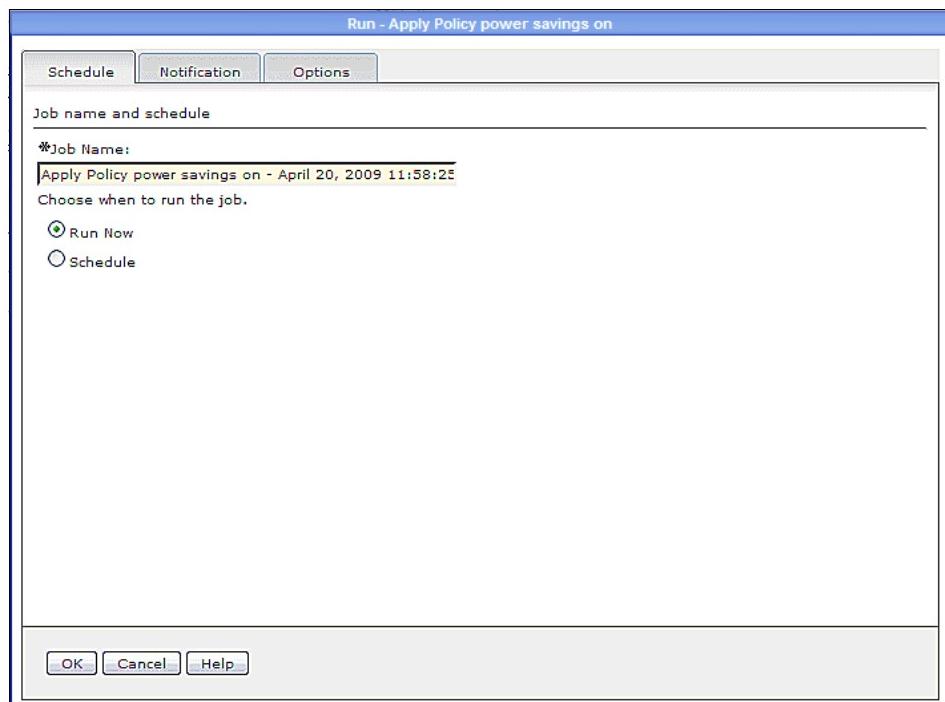


Figure 10-39 Run now - policy apply option

You also have the option of scheduling when to run the power policy. This feature is used to apply a power policy unattended. This would be useful for setting up automatic application of a policy to turn on/off power savings for

example. Figure 10-40 on page 429 provides an example of the settings to schedule a policy.

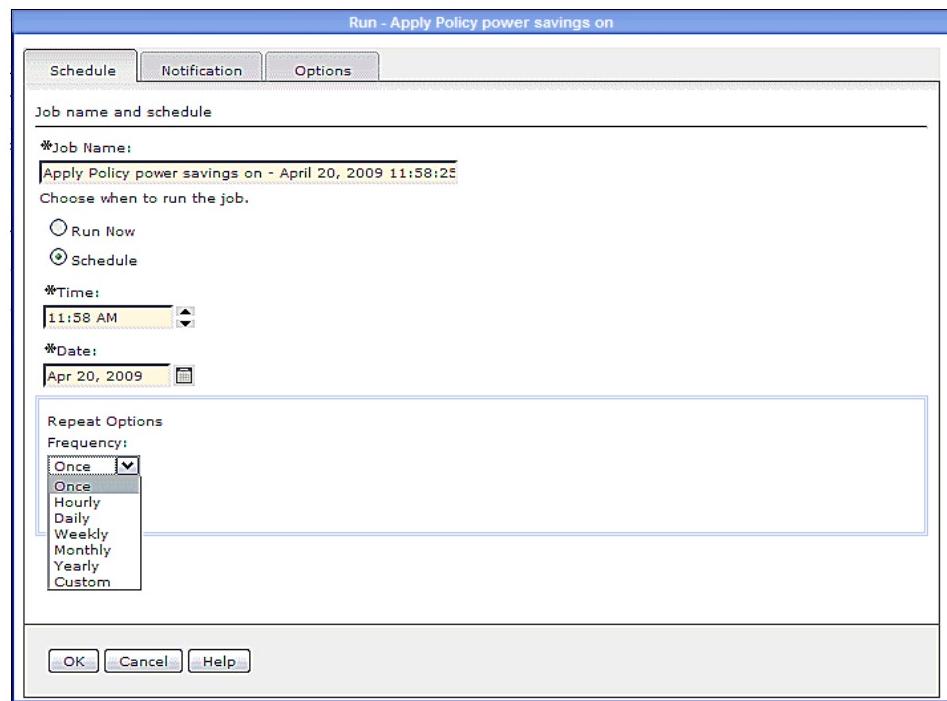


Figure 10-40 Policy schedule options

You can also set the system to send you an E-mail when the policy is applied. Modify the **Notification** tab settings for your correct contact information. Figure 10-41 on page 430 shows an example of the Notification tab.

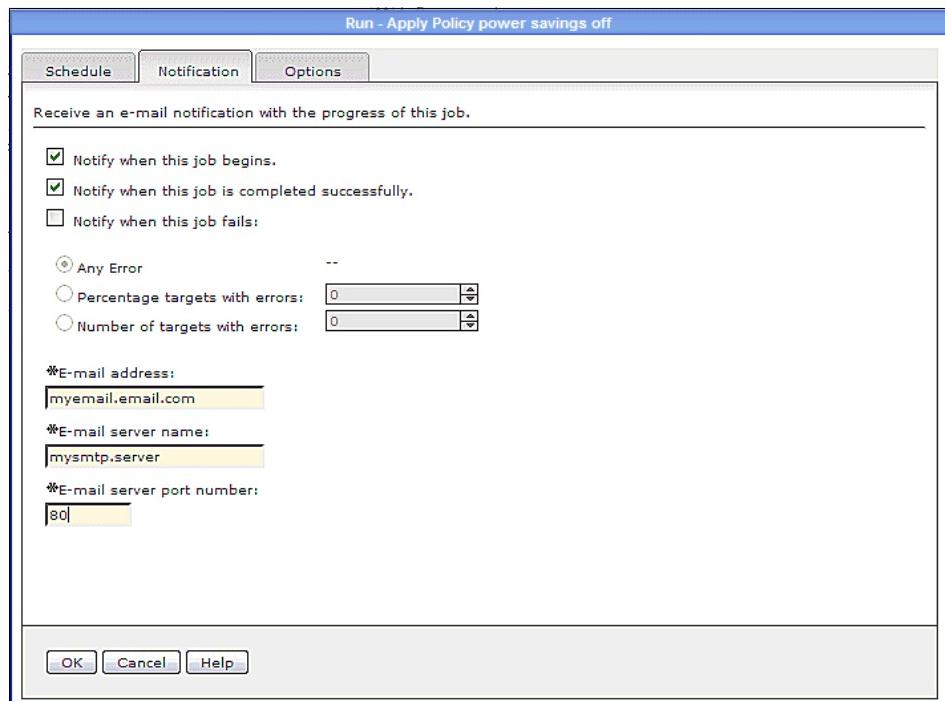


Figure 10-41 Notification tab

The Options tab will allow you to set which time base to use, either management server or local system time. You also have the option to allow the policy action to fail if the system is not available or run when the system becomes available.

Figure 10-42 on page 431 shows an example of these settings.

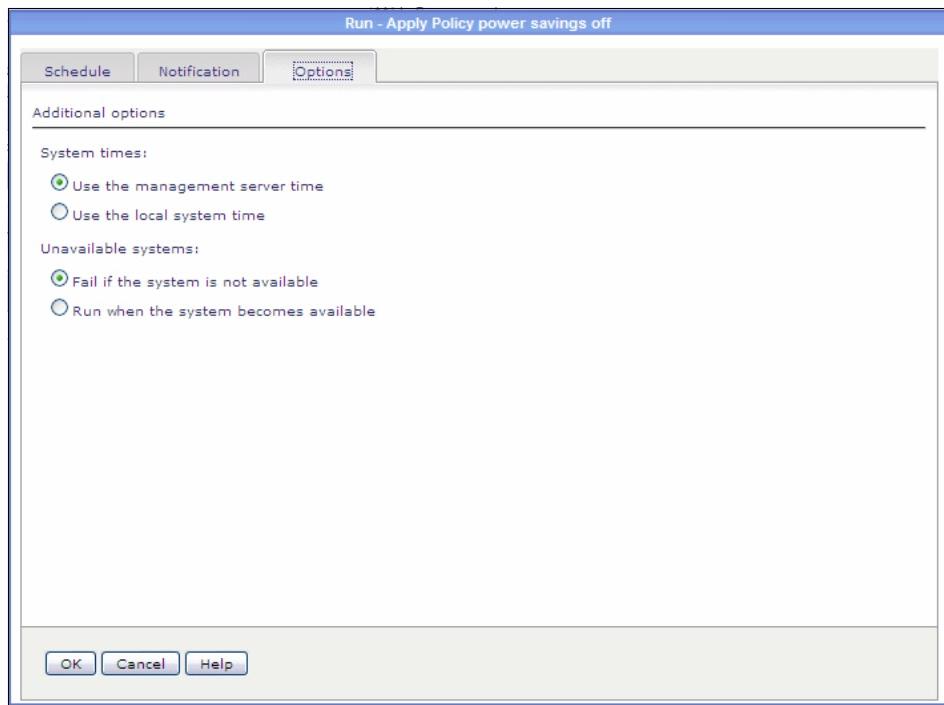


Figure 10-42 Policy options tab

Active Energy Manager can also be controlled through the command line interface. Many of the CLIs are useful to the IBM BladeCenter management.

Information about the smcli interface can be found here:

http://publib.boulder.ibm.com/infocenter/systems/index.jsp?topic=/aem_410/frb0_main.html

Information about IBM Systems Director command line interface can be found here:

http://publib.boulder.ibm.com/infocenter/systems/index.jsp?topic=/director.cli_6.1/fqm0_r_cli_smcli.html

For more information on IBM Systems Director or Active Energy Manager be sure to reference the links at the beginning of this chapter.



Performing Live Partition Mobility

This chapter discusses the requirements and configuration procedures to perform Live Partition Mobility between a IBM BladeCenter JS23 and JS43 blades.

We cover the following in this chapter:

- ▶ “Requirements” on page 434
- ▶ “Preparation” on page 438
- ▶ “Migrating the LPAR” on page 448

Additional information on Live Partition Mobility architecture, mechanisms and advanced topics can be found in the Redbook *IBM PowerVM Live Partition Mobility*, SG24-7460.

For the most current information, documentation, known problems, workarounds etc. go to the Live Partition Mobility support web-site at this locations:

<http://www14.software.ibm.com/webapp/set2/sas/f/pm/home.html>

11.1 Requirements

Partition mobility places certain demands on hardware, software, network and storage configurations. These considerations need to be reviewed early in the setup of an IBM BladeCenter JS23 or JS43 to avoid reconfiguration and rework.

11.1.1 Hardware

The IBM BladeCenter JS23 or JS43 requires a Fibre Channel HBA expansion card for SAN connectivity. All storage assigned to a logical partition must have external backing devices that are visible to both the local and remote VIOS systems.

11.1.2 Firmware

In a new blade environment the recommendation would be to have the most current system firmware available installed. For the hardware configuration used in the writing of this book the source and target blades had two different levels, the current and previous levels of firmware. This configuration was chosen to represent a customer environment where Live Partition Mobility may be used for system maintenance reasons. IBM BladeCenter JS23 and JS43 firmware updates can be downloaded from

<http://www.ibm.com>

Select the **Support & Downloads** link, then **BladeCenter** from the drop-down box to be directed to the Support for IBM BladeCenter page. Select **BladeCenterJS23** or **BladeCenterJS43** from the Product family drop-down box, then click **Go** to display a list of related blade and BladeCenter updates. Blade system firmware will have a prerequisite AMM firmware that will be identified in the firmware readme file.

11.1.3 VIOS version

Similarly to system firmware, the VIOS version and fixpack level should be the most recent. To display the current code level from the UI, click **Updates** from the navigation area. The Management Partition Updates view will open and the code level shown in Figure 11-1 on page 435. If the workstation that you are using has Internet access, the link displayed on the Management Partition Updates page will take you to a download site for newer updates and fixes, if available. The link to the Virtual I/O Server is also available here:

<http://techsupport.services.ibm.com/server/vios/download>



Figure 11-1 Management Partition Updates view

From the CLI use the **ioslevel** command to display the VIOS version and fixpack level, as shown in Example 11-1. In this example the VIOS version is 2.1.1.0 and has not had any fixpacks installed.

Example 11-1 ioslevel command

```
$ ioslevel  
2.1.1.0
```

An example of a previous release with a fixpack installed is shown in Example 11-2.

Example 11-2 ioslevel command showing fixpack installed

```
$ ioslevel  
2.1.0.10-FP-20.1
```

11.1.4 PowerVM Enterprise

PowerVM Enterprise Edition is an optional feature on an IBM BladeCenter JS23 or JS43 and is required to enable Partition Mobility. To determine if this capability is available use the **lssyscfg** command. Example 11-3 shows the lssyscfg returning a value of 1 to indicate active or live partition mobility capability.

Example 11-3 lsset command

```
$ lssyscfg -r sys -F active_lpar_mobility_capable  
1
```

If Partition Mobility is not enabled and the feature was purchased with the blade, the activation key should be available on the IBM Capacity on Demand (CoD) Web site at this link:

<http://www-912.ibm.com/pod/pod>

Enter the system type and serial number on the CoD site and submit. A list of available activation codes or keys with a type and description will be displayed. If PowerVM Enterprise Edition was not purchased with the IBM BladeCenter JS23 or JS43, it can be upgraded through the Miscellaneous Equipment Specification (MES) process.

Entering an enablement key through the IVM UI

The PowerVM key is entered from the UI by clicking the **Enter PowerVM Edition Key** link in the navigation area. The window that opens will have a box to enter the key number. Type in the key and select **Apply** as shown in Figure 11-2 on page 437. When PowerVM Enterprise is enabled, a Mobility section is added to the More Tasks drop-down box on the View/Modify Partitions view.

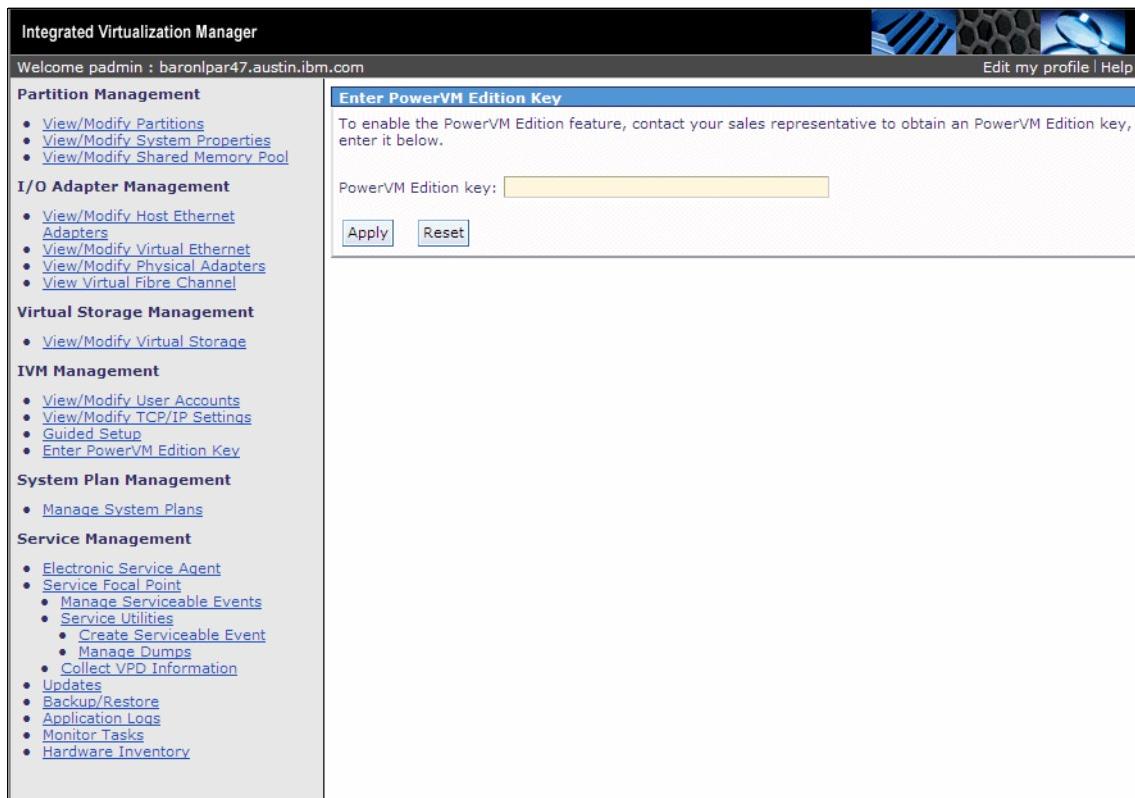


Figure 11-2 PowerVM Enterprise key entry

11.1.5 LPAR OS versions

The running operating system in the mobile partition must be AIX or Linux. The currently supported operating systems for Live Partition Mobility are:

- ▶ AIX 5L V5.3 with 5300-07 Technology Level or later
- ▶ AIX V6.1 or later
- ▶ Red Hat Enterprise Linux Version 5.1 or later
- ▶ SUSE Linux Enterprise Services 10 (SLES 10) Service Pack 1 or later

11.2 Preparation

This section describes the settings and configurations that must be verified and possibly changed to prepare the local and remote VIOS systems and partitions for partition mobility.

11.2.1 VIOS (source and target) requirements

We'll start with VIOS (source and target) considerations.

Memory region size

The memory region size is the smallest block of memory that can be assigned to or changed in an LPAR. The current setting can be reviewed by clicking the **View/Modify System Properties** link from the Navigation area and then selecting the **Memory** tab. By default is dependent on the amount of system memory installed. The value is set to 32MB (automatic) as shown in Figure 11-3 on page 439. The memory region size must be the same for the source and target VIOS systems. If the value is changed, a VIOS power down and restart is required (not just a basic shutdown reboot).

Figure 11-3 Memory region size

Storage and hdisk reserve policy

Only physical volumes (LUNs) visible to the VIOS as a hdisk assigned to an LPAR can be used in mobile partitions. The same physical volumes must also be visible to both the local and remote VIOS systems. The reserve policy of the hdisk must be changed from the default single_path to no_reserve. The reserve policy is changed on an hdisk from both VIOS systems. The `lsdev -dev hdisk# -attr | grep reserve` command can be used to review the current policy setting. The `chdev` command is used to change the attribute of the hdisk. Example 11-4 shows hdisk1 with the default reserve policy of single_path, the `chdev` command being issued, and the new reserve policy of no_reserve.

Example 11-4 Changing a hdisk reserve policy

```
$ lsdev -dev hdisk1 -attr | grep reserve
reserve_policy single_path             Reserve Policy
True
```

```
$ chdev -dev hdisk1 -attr reserve_policy=no_reserve  
hdisk1 changed
```

```
$ lsdev -dev hdisk1 -attr | grep reserve  
reserve_policy no_reserve Reserve Policy  
True
```

Note: The reserve policy cannot be changed on the source VIOS when the disks are assigned to an LPAR. The command will fail with the following message:

Some error messages may contain invalid information for the Virtual I/O Server environment.

Method error (/etc/methods/chgfcparray):

0514-062 Cannot perform the requested function because the specified device is busy.

If the reserve policy has not been changed, it can quickly be identified by looking on the target VIOS. From the target system IVM UI Navigation area click the **View/Modify Virtual Storage** link. In the View/Modify Virtual Storage view click the **Physical Volumes** tab and review the size column. If any of the sizes are unknown, it is likely the reserve policy has not been changed and the validation process will fail. Figure 11-4 on page 441 shows hdisk0 and hdisk7 in the unknown size condition.

Select	Name	Storage Pool	Assigned Partition	Size	Physical Location Code
<input type="checkbox"/>	hdisk0			Unknown	U78A5.001.WIH01B7-P1-C7-T2-W203300A0B811A662
<input type="checkbox"/>	hdisk1	rootvg (Default)		68.37 GB	U78A5.001.WIH01B7-P1-D1
<input type="checkbox"/>	hdisk2			10 GB	U78A5.001.WIH01B7-P1-C7-T2-W203300A0B811A662
<input type="checkbox"/>	hdisk3	Phobes - RHEL5-U2 (2)		5 GB	U78A5.001.WIH01B7-P1-C7-T2-W203300A0B811A662
<input type="checkbox"/>	hdisk4			5 GB	U78A5.001.WIH01B7-P1-C7-T2-W203300A0B811A662
<input type="checkbox"/>	hdisk5			5 GB	U78A5.001.WIH01B7-P1-C7-T2-W203300A0B811A662
<input type="checkbox"/>	hdisk6			50 MB	U78A5.001.WIH01B7-P1-C7-T2-W203300A0B811A662
<input type="checkbox"/>	hdisk7			Unknown	U78A5.001.WIH01B7-P1-C7-T2-W203300A0B811A662

Figure 11-4 hdisk reserve policy not set correctly

When the validation process is run, an error message similar to Figure 11-5 on page 442 will be displayed.

This problem can be resolved by performing the following steps:

1. Shutting down the mobile LPAR on the local VIOS if running.
2. Modifying the mobile LPAR hdisk assignments on the local VIOS to *none*.
3. Using the **chdev** command to change the hdisks reserve policy to *no_reserve*.
4. Modifying the mobile LPAR hdisk assignments to the original assignments.
5. Refreshing the View/Modify Virtual Storage view on the remote VIOS.

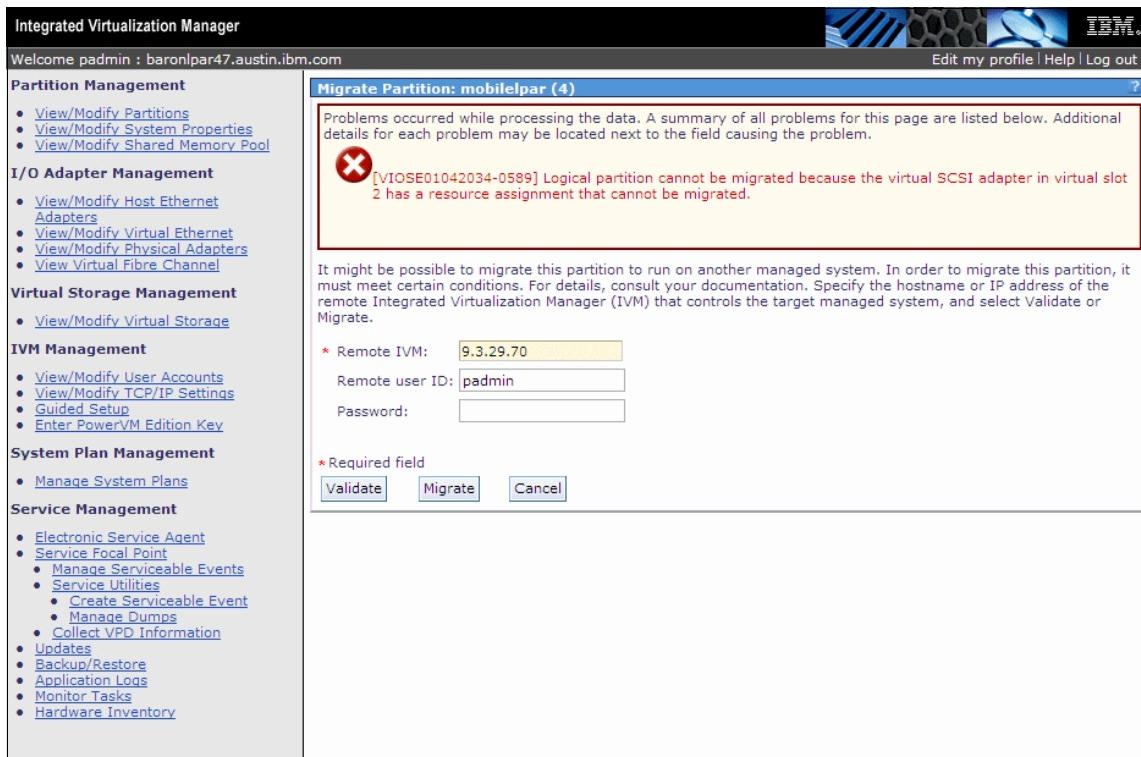


Figure 11-5 Partition Migration validation error message for target storage

11.2.2 Networking

The mobile LPAR external network communication must be through a Shared Ethernet Adapter (SEA). The use of logical ports on a Host Ethernet Adapter (HEA) or physical adapters assigned to the LPAR cannot be used and must be removed if assigned. SEA adapter creation is covered in 4.5.2, “Virtual Ethernet Adapters and SEA” on page 103.

The Resource Monitoring and Control (RMC) daemon must be active on the mobile partition or the validate and migration process will fail. The IVM UI can be used to verify the status as described in 4.7.11, “Partition properties changes and DLPAR operations” on page 156.

The CLI **lssyscfg** command can also be used to determine the RMC status as shown in Example 11-5.

Example 11-5 lssyscfg command to determine the RMC status

```
$ lssyscfg -r lpar -F name,rmc_state
```

VIOS-Neptune,active
Phobes - RHEL5-U2,inactive
Mars - AIX 6.1,active

Note: Linux partitions must have the Dynamic Reconfiguration Tools package for HMC- or IVM-managed servers installed from the Service and Productivity tools Web site at:

<https://www14.software.ibm.com/webapp/set2/sas/f/lopdiags/home.html>

Service and Productivity tools are discussed in Appendix D, “Service and productivity tools for Linux” on page 545.

11.2.3 Partition requirements

In this section we discuss the partition requirements for performing Live Partition Mobility.

Dedicated or Shared Memory

For a partition that is using dedicated memory, the target VIO Server must have adequate available memory to contain the moving partition. |

Note: In cases where the available memory on the target VIO Server and the dedicated memory configured in the mobile partition match it will fail the validation process. The creation of an additional LPAR on the target VIO Server will cause firmware to reserve additional memory.

A partition that is using shared memory on the source VIO Server must have a shared memory pool defined on the target VIO Server with adequate available resources for the mobile partition.

Processor Compatibility mode

The processor compatibility mode has two settings: the current value and the preferred value for the logical partition and can be displayed from the Partition Properties Processor tab.

The Current value indicates the negotiated compatibility mode for the logical partition. This is the value that the logical partition is currently using.

The following values are possible for the current setting of the processor compatibility mode.

POWER6

This mode is possible for both POWER6 and POWER6 + processor based servers. This mode indicates that the operating environment for the partition is using all the standard capabilities of the POWER6 processor.

POWER6+

This mode is possible for POWER6 + processor based servers. This mode indicates that the operating environment for the partition is using all the standard capabilities of the POWER6 + processor.

POWER6 Enhanced

This mode is possible for POWER6 processor based servers. This mode indicates that the operating environment for the partition is using all the standard features of the POWER6 processor and also supports using additional floating-point instructions in the applications that use the POWER6 processor.

POWER6+ Enhanced

This mode is possible for POWER6 + processor based servers. This mode indicates that the operating environment for the partition is using all the standard features of the POWER6 processor and also supports using additional floating-point instructions in the applications that use the POWER6 + processor.

After you select the preferred processor compatibility mode, you must shut down and restart the logical partition so that the hypervisor can check the preferred mode and the operating environment to change the value for the current processor compatibility mode.

For active logical partition migration, the destination server must support both preferred and current processor compatibility modes of the mobile partition. For inactive migrations, the destination server must support only the preferred processor compatibility mode of the mobile partition.

The CPU type on JS23/JS43 is Power6+ and has different compatibility modes that can be selected as shown in Figure 11-6 on page 445.

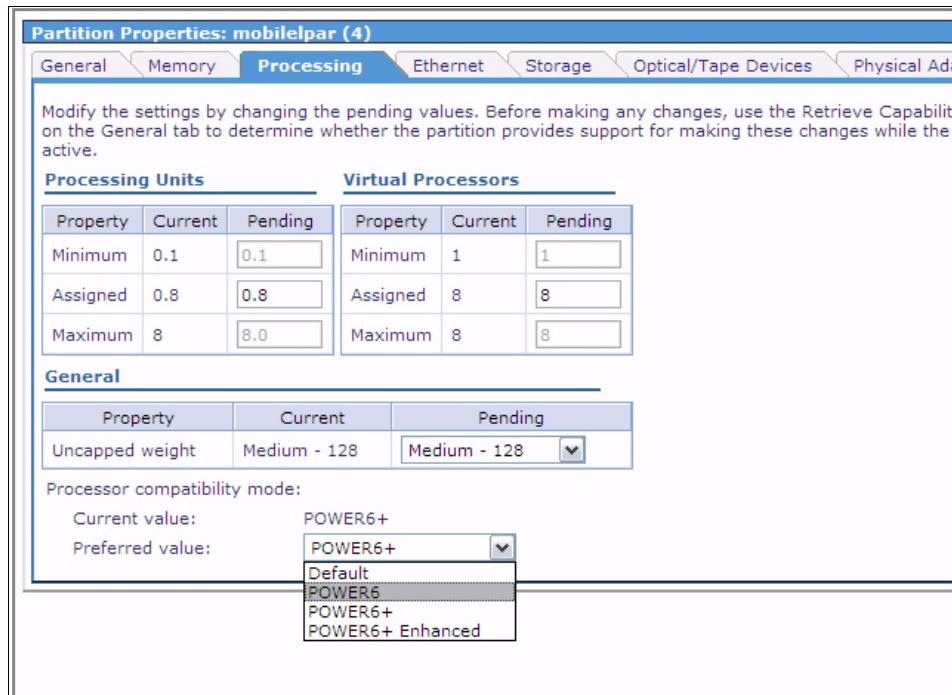


Figure 11-6 Processor compatibility mode on JS23/JS43

JS12 and JS22 blades used POWER6 technology and can be configured for the processor compatibility modes as shown in Figure 11-7 on page 446.

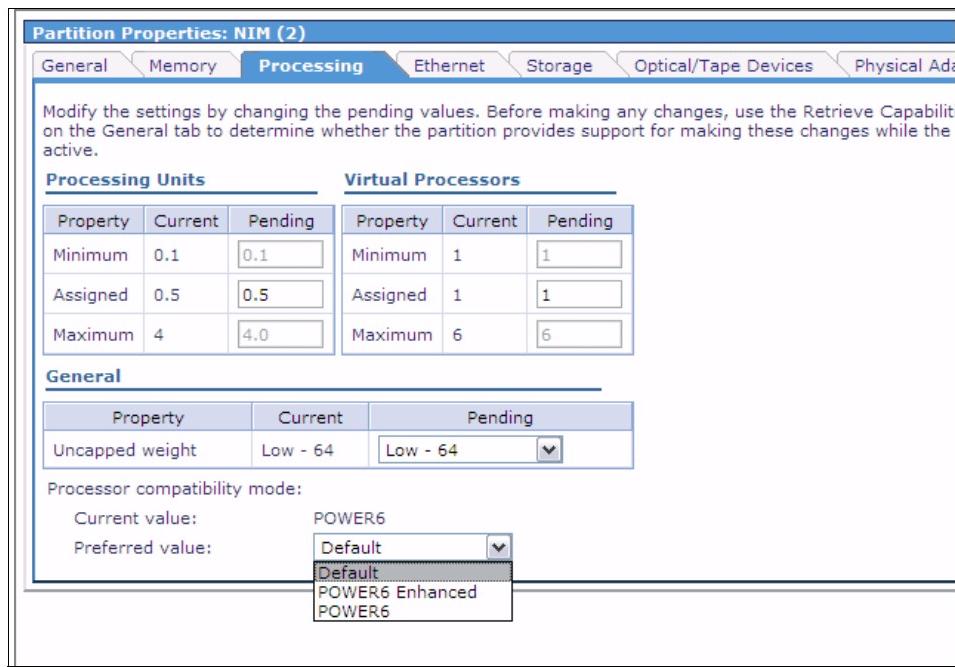


Figure 11-7 Processor compatibility mode on JS12/JS22

The requirement is that the source and target blades have the ability to match processor compatibility modes. Currently for POWER6 based blades the only common processor compatibility mode is POWER6. An LPAR running in POWER6 mode on a JS12 could migrate to a JS23 or JS43. If the JS12 LPAR was running in POWER6 Enhanced migration to a JS23 or JS43 would not be possible without a mode change first on the JS12 to POWER6 mode.

To perform Live Partition Mobility from IBM BladeCenter JS23/JS43 blades to IBM BladeCenter JS12/JS22 the processor compatibility mode on JS23/JS43 must be changed to POWER6 as shown in Figure 11-8 on page 447.

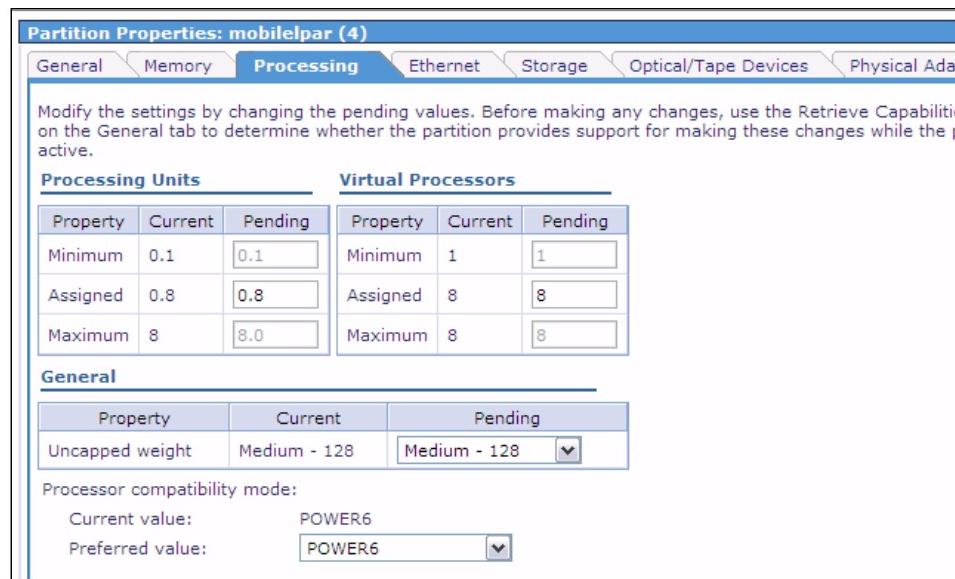


Figure 11-8 Change the processor compatibility mode on JS23/JS43

Virtual optical devices

All virtual optical devices must be removed from the mobile partition before a successful validation and migration can occur. The example shown in Figure 11-9 on page 448 indicates that the virtual device vtopt0 is still assigned to the mobile partition. The device can be removed by unchecking the box and clicking **OK**.

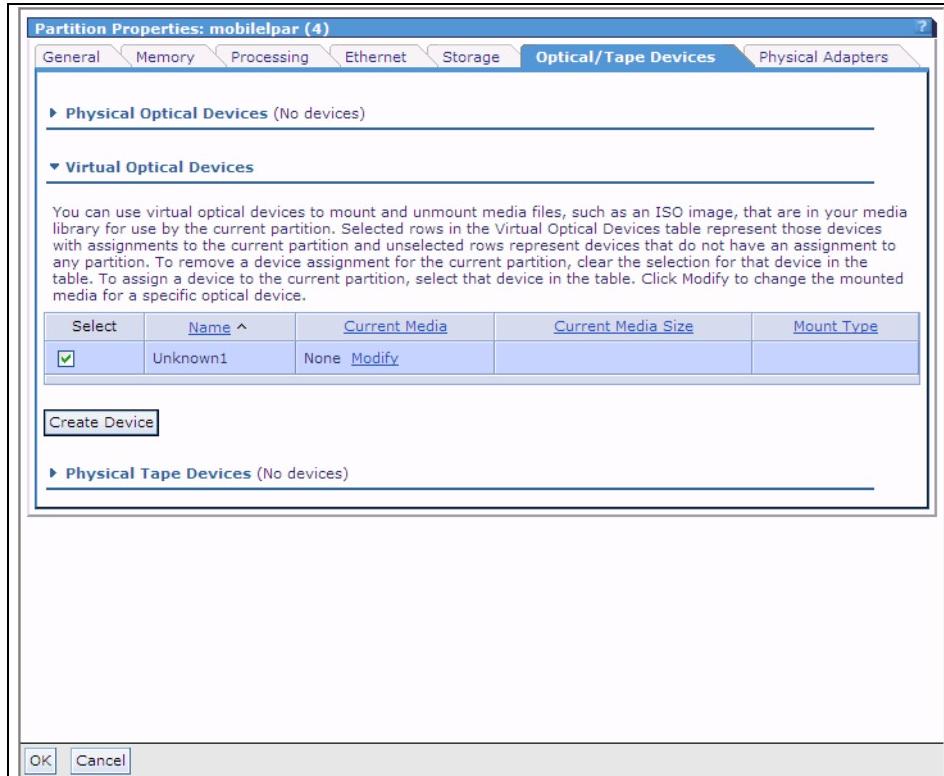


Figure 11-9 Virtual optical device to be removed

11.3 Migrating the LPAR

The following sections describe how to use the IVM UI and CLI to validate, migrate, and check status on mobile LPAR.

11.3.1 Using the IVM UI

Let us first see how we can perform an LPAR migration with IVM.

Validate

The migration process is started by first selecting **View/Modify Partitions** from the Navigation area. With this view open, use the check box for the desired mobile partition and select **Migrate** from the More Tasks drop-down box as shown in Figure 11-10 on page 449.

The screenshot shows the IVM interface with the following details:

- Partition Management** section on the left:
 - [View/Modify Partitions](#)
 - [View/Modify System Properties](#)
 - [View/Modify Shared Memory Pool](#)
- I/O Adapter Management** section:
 - [View/Modify Host Ethernet Adapters](#)
 - [View/Modify Virtual Ethernet](#)
 - [View/Modify Physical Adapters](#)
 - [View Virtual Fibre Channel](#)
- Virtual Storage Management** section:
 - [View/Modify Virtual Storage](#)
- IVM Management** section:
 - [View/Modify User Accounts](#)
 - [View/Modify TCP/IP Settings](#)
 - [Guided Setup](#)
 - [Enter PowerVM Edition Key](#)
- System Plan Management** section:
 - [Manage System Plans](#)
- Service Management** section:
 - [Electronic Service Agent](#)
 - [Service Focal Point](#)
 - [Manage Serviceable Events](#)
 - [Service Utilities](#)
 - [Create Serviceable Event](#)
 - [Manage Dumps](#)
 - [Collect VPD Information](#)
 - [Updates](#)
 - [Backup/Restore](#)
 - [Application Logs](#)
 - [Monitor Tasks](#)
 - [Hardware Inventory](#)

View/Modify Partitions section on the right:

- To perform an action on a partition, first select the partition or partitions, and then select the task.

System Overview table:

Total system memory:	8 GB	Total processing units:	8
Memory available:	3.28 GB	Processing units available:	8
Reserved firmware memory:	640 MB	Processor pool utilization:	0%
System attention LED:	Inactive		

Partition Details table:

Select	ID	Name	State	Uptime	Memory	Proc	... More Tasks ...		
<input type="checkbox"/>	1	is43-vios	Running	17.74 Hours	1.5 GB	8	... More Tasks ... Open terminal window Delete Create based on Operator panel service functions Reference Codes		
<input type="checkbox"/>	2	rhel53	Not Activated		1 GB	1	... Mobility ... Migrate Status <hr/> Properties		
<input type="checkbox"/>	3	IBMi	Not Activated		1 GB	1			
<input checked="" type="checkbox"/>	4	mobilepar	Running	10.7 Minutes	608 MB	8	0.8	0.02	

Figure 11-10 Partition Migrate option

The Migrate Partition view will open with the mobile partition name appended to the window name. Enter the remote or target IVM-controlled system IP address, remote user ID and password as shown in Figure 11-11 on page 450. Click **Validate** to start the validation process.

Note: The Partition Migration view requests the Remote IVM or HMC IP address. At the time of this publication, IVM to HMC migrations are not supported.

Integrated Virtualization Manager

Welcome padmin : baronlpar47.austin.ibm.com

Edit my profile

Partition Management

- [View/Modify Partitions](#)
- [View/Modify System Properties](#)
- [View/Modify Shared Memory Pool](#)

I/O Adapter Management

- [View/Modify Host Ethernet Adapters](#)
- [View/Modify Virtual Ethernet](#)
- [View/Modify Physical Adapters](#)
- [View Virtual Fibre Channel](#)

Virtual Storage Management

- [View/Modify Virtual Storage](#)

IVM Management

- [View/Modify User Accounts](#)
- [View/Modify TCP/IP Settings](#)
- [Guided Setup](#)
- [Enter PowerVM Edition Key](#)

System Plan Management

- [Manage System Plans](#)

Service Management

- [Electronic Service Agent](#)
- [Service Focal Point](#)
- [Manage Serviceable Events](#)
- [Service Utilities](#)
 - [Create Serviceable Event](#)
 - [Manage Dumps](#)
 - [Collect VPD Information](#)
- [Updates](#)
- [Backup/Restore](#)
- [Application Logs](#)
- [Monitor Tasks](#)
- [Hardware Inventory](#)

Migrate Partition: mobilelpar (4)

It might be possible to migrate this partition to run on another managed system. In order to migrate, certain conditions must be met. For details, consult your documentation. Specify the hostname or IP address of the remote Integrated Virtualization Manager (IVM) that controls the target managed system, and select Migrate.

* Remote IVM:

Remote user ID:

Password:

* Required field

Validate **Migrate** **Cancel**

Figure 11-11 Partition Mobility validation

At the end of the successful validation process, the Migrate Partition window will be updated similar to Figure 11-12.

The screenshot shows the Integrated Virtualization Manager (IVM) interface. On the left, there is a sidebar with various management categories and their sub-links:

- Partition Management**: View/Modify Partitions, View/Modify System Properties, View/Modify Shared Memory Pool
- I/O Adapter Management**: View/Modify Host Ethernet Adapters, View/Modify Virtual Ethernet, View/Modify Physical Adapters, View Virtual Fibre Channel
- Virtual Storage Management**: View/Modify Virtual Storage
- IVM Management**: View/Modify User Accounts, View/Modify TCP/IP Settings, Guided Setup, Enter PowerVM Edition Key
- System Plan Management**: Manage System Plans
- Service Management**: Electronic Service Agent, Service Focal Point (Manage Serviceable Events, Service Utilities, Create Serviceable Event, Manage Dumps, Collect VPD Information), Updates, Backup/Restore, Application Logs, Monitor Tasks, Hardware Inventory

The main content area is titled "Migrate Partition: mobilelpar (4)". It displays a message: "The operation completed successfully." Below this, a note states: "It might be possible to migrate this partition to run on another managed system. In order to migrate, the target system must meet certain conditions. For details, consult your documentation. Specify the hostname or IP address of the remote Integrated Virtualization Manager (IVM) that controls the target managed system, and select the Migrate button." There are three input fields: "Remote IVM:" (containing "9.3.29.70"), "Remote user ID:" (containing "padmin"), and "Password:" (containing "*****"). A note below the password field says "* Required field". At the bottom are three buttons: Validate, Migrate (which is highlighted in blue), and Cancel.

Figure 11-12 Partition Migration validation success

Figure 11-13 shows the results of the validation process that discovered a problem that would prevent a migration. This error message was generated because of a virtual SCSI assignment that could not be migrated. In this example the problem was due to a virtual optical device that had an assignment to the mobile partition. Another example is shown in Figure 11-4 on page 441, where the validation process could not find the required storage on the remote system. The validation process must complete without errors before a partition migration can be attempted.

The screenshot shows the IVM web interface with the following details:

- Header:** Integrated Virtualization Manager, Welcome padmin : baronpar47.austin.ibm.com, Edit my profile | Help | Log out
- Left Sidebar (Navigation):**
 - Partition Management:** View/Modify Partitions, View/Modify System Properties, View/Modify Shared Memory Pool
 - I/O Adapter Management:** View/Modify Host Ethernet Adapters, View/Modify Virtual Ethernet, View/Modify Physical Adapters, View Virtual Fibre Channel
 - Virtual Storage Management:** View/Modify Virtual Storage
 - IVM Management:** View/Modify User Accounts, View/Modify TCP/IP Settings, Guided Setup, Enter PowerVM Edition Key
 - System Plan Management:** Manage System Plans
 - Service Management:** Electronic Service Agent, Service Focal Point (Manage Serviceable Events, Service Utilities: Create Serviceable Event, Manage Dumps, Collect VPD Information), Updates, Backup/Restore, Application Logs, Monitor Tasks, Hardware Inventory
- Middle Content Area:**

Migrate Partition: mobilepar (4)

Problems occurred while processing the data. A summary of all problems for this page are listed below. Additional details for each problem may be located next to the field causing the problem.

Error Message: [VIOSE01042034-0589] Logical partition cannot be migrated because the virtual SCSI adapter in virtual slot 2 has a resource assignment that cannot be migrated.

It might be possible to migrate this partition to run on another managed system. In order to migrate this partition, it must meet certain conditions. For details, consult your documentation. Specify the hostname or IP address of the remote Integrated Virtualization Manager (IVM) that controls the target managed system, and select Validate or Migrate.

* Remote IVM: (highlighted)

Remote user ID:

Password:

* Required field

Buttons: Validate, Migrate, Cancel

Figure 11-13 Failed validation because of improper virtual SCSI device assignment

Migrate

With a successful completion of the validation process the migrate step can be started. Click **Migrate** to begin the migration process. As part of the migration process, a validate is run again and at the end of this step a Migrate Status view will display, as shown in Figure 11-14.

The screenshot shows the Integrated Virtualization Manager (IVM) web interface. On the left, there is a navigation sidebar with several management sections:

- Partition Management**: View/Modify Partitions, View/Modify System Properties, View/Modify Shared Memory Pool
- I/O Adapter Management**: View/Modify Host Ethernet Adapters, View/Modify Virtual Ethernet, View/Modify Physical Adapters, View Virtual Fibre Channel
- Virtual Storage Management**: View/Modify Virtual Storage
- IVM Management**: View/Modify User Accounts, View/Modify TCP/IP Settings, Guided Setup, Enter PowerVM Edition Key
- System Plan Management**: Manage System Plans
- Service Management**: Electronic Service Agent, Service Focal Point, Manage Serviceable Events, Service Utilities (Create Serviceable Event, Manage Dumps), Collect VPD Information, Updates, Backup/Restore, Application Logs, Monitor Tasks, Hardware Inventory

The main content area is titled "Migrate Status". It contains a message: "The following partitions are currently migrating. You may stop this operation or continue to the next step. If the status shows an error, you should select Recover, which will attempt to complete the migration." Below this is a table titled "Partitions Migrating From This System".

Select	Partition	Migration Status	Percent Complete	Remote Platform Manager
<input type="checkbox"/>	mobilelpar (4)	Migration Starting	45%	9.3.29.70

At the bottom of the table are "OK" and "Cancel" buttons.

Figure 11-14 Migrate Status view

The Migrate Status view can be accessed directly from the View/Modify Partitions window. Check the mobile partition box, then select **Status** under the Mobility section of the More Tasks drop-down box as shown in Figure 11-15. Also note in this same figure that the state of the mobile partition has changed from *Running* to *Migrating- Running*.

Select	ID	Name	State	Uptime	Memory	Processor
<input type="checkbox"/>	1	js43-vios	Running	18.07 Hours	1.5 GB	8
<input type="checkbox"/>	2	rhel53	Not Activated		1 GB	1
<input type="checkbox"/>	3	IBMi	Not Activated		1 GB	1
<input type="checkbox"/>	6	mobilepar	Migrating - Running	30.1 Minutes	608 MB	8

Figure 11-15 Starting the status view for a mobile partition

Figure 11-16 shows the View/Modify Partitions view on the remote IVM, indicating migration has started.

Note: The mobile partition will retain the same LPAR ID number if available on the remote system, otherwise it will be assigned the first available ID number.

The screenshot shows the Integrated Virtualization Manager (IVM) interface. The left sidebar contains a navigation menu with several sections: Partition Management, I/O Adapter Management, Virtual Storage Management, IVM Management, System Plan Management, and Service Management. The main content area is titled "View/Modify Partitions". It includes a "System Overview" section with system statistics and a "Partition Details" section where partitions are listed in a table. The "mobilelpar" partition is highlighted in yellow and is shown with a status of "Migrating - Running".

Select	ID ^	Name	State	Uptime	Memory	Processors	Entitled Processing Units
<input type="checkbox"/>	1	is23-vios	Running	3.11 Hours	2 GB	4	0.4
<input type="checkbox"/>	2	IBMI_2	Not Activated		256 MB	1	0.5
<input type="checkbox"/>	3	JS23AMSlpar3	Not Activated		4 GB	4	0.4
<input type="checkbox"/>	4	JS23DMLpar4	Not Activated		4 GB	4	0.4
<input type="checkbox"/>	5	JS23DPlpar5	Not Activated		1 GB	1	1.0
<input type="checkbox"/>	6	mobilelpar	Migrating - Running		608 MB	8	0.8

Figure 11-16 Remote IVM indicating migration in progress

At the end of the migration process the State of the mobile partition changes from *Migrating - Running* to *Running* as shown in Figure 11-17 on the formerly remote system. On the original local system the mobile partition is removed from the View/Modify Partition view.

Select	ID ^	Name	State	Uptime	Memory	Processors	Entitled Processing Units
<input type="checkbox"/>	1	is23-vios	Running	3.14 Hours	2 GB	4	0.4
<input type="checkbox"/>	2	IBMI_2	Not Activated		256 MB	1	0.5
<input type="checkbox"/>	3	JS23AMSIpar3	Not Activated		4 GB	4	0.4
<input type="checkbox"/>	4	JS23DMP1par4	Not Activated		4 GB	4	0.4
<input type="checkbox"/>	5	JS23DPIpar5	Not Activated		1 GB	1	1.0
<input type="checkbox"/>	6	mobile1par	Running	32.1 Minutes	608 MB	8	0.8

Figure 11-17 Partition migration complete to remote system

11.3.2 From the command line

The IVM **migr1par** command is used to validate and migrate the mobile partition from one IVM-managed system to another. Example 11-6 shows the **validate** command and the result when a virtual optical device is still assigned to the mobile partition. A successful validation process will return to the \$ prompt with no other messages.

Example 11-6 *migr1par command validate option*

```
$ migr1par -o v -t Server-7998-61X-SN7157008 --ip 172.16.1.100 --id 5
```

[VIOSE01042034-0418] The partition cannot be migrated because the virtual SCSI server adapter has a resource assignment that cannot be migrated.

The -o flag or operation has the following options:

- ▶ s - stop
- ▶ m - validate and migrate
- ▶ r - recover
- ▶ v - validate

The -t flag in Example 11-6 on page 456 specifies the remote managed system. The -t flag requires a system name and IP address.

Note: The system name is not the same as the host name. The system name can be obtained from the **lssyscfg** command as follows:

```
$ lssyscfg -r sys -F name  
Server-7998-61X-SN7157008
```

The system name can also be viewed or easily changed from the IVM UI by clicking **View/Modify System Properties** in the Navigation area.

The --id flag specifies the mobile partition to be validated or migrated.

A successful validation will have a zero return code.

Example 11-7 shows the **migr1par** command with the operation flag with the m option for validate and migrate.

Example 11-7 migr1par command validate and migrate option

```
$ migr1par -o m -t Server-7998-61X-SN7157008 --ip 172.16.1.100 --id 5  
$
```

The status of the mobile partition can be monitored by use of the **lssyscfg** command as shown in Example 11-8 or the **ls1parmigr** command used in Example 11-9 on page 458 for the LPAR named Mars. Repeating the **ls1parmigr** command will show a change in the bytes transmitted and the bytes remaining as the migration progresses.

Example 11-8 lssyscfg command used to check migrating partition status

```
$ migr1par -o m -t Server-7998-61X-SN7157008 --ip 172.16.1.100 --id 5 &
```

```
[1]      24076366  
  
$ lssyscfg -r lpar -F name,state  
VIOS-Neptune,Running  
Phobes - RHEL5-U2,Running  
Mars - AIX 6.1,Migrating - Running
```

Example 11-9 lsparmigr command used to check migrating partition status

```
$ migrlpar -o m -t Server-7998-61X-SN7157008 --ip 172.16.1.100 --id 5 &  
[1]      24228082  
  
$ lsparmigr -r lpar  
name=VIOS-Neptune,lpard_id=1,migration_state=Not Migrating  
name=Phobes - RHEL5-U2,lpard_id=2,migration_state=Not Migrating  
name=Mars - AIX 6.1,lpard_id=5,migration_state=Migration In  
Progress,migration_type=active,dest_sys_name=Server-7998-61X-SN7157008,  
dest_lpar_id=5,source_msp_name=VIOS-Neptune,source_msp_id=1,dest_msp_na  
me=Jupiter,dest_msp_id=1,bytes_transmitted=1117792957,bytes_remaining=1  
7162240,remote_manager=172.16.1.100,remote_user=padmin
```

Once the migration is complete, the instance of the mobile LPAR will no longer appear on the original source VIOS.



System maintenance and diagnostics

This chapter discuss methods and best practices related to some important IBM BladeCenter JS23 and JS43 Express maintenance topics, such as:

- ▶ “Firmware updates” on page 460.
- ▶ “System diagnostics” on page 472

12.1 Firmware updates

IBM periodically makes firmware updates available for you to install on the IBM BladeCenter JS23 and JS43 Express, the management module, or expansion cards in the blade server.

IBM BladeCenter JS23 and JS43 Express have a large firmware image, making it impossible to perform firmware updates through the Advanced Management Module. For these blade servers, one of the following process should be chosen to update system's firmware:

- ▶ In-band operating system capabilities, such as the `update_flash` command for Linux and AIX, or the command for Virtual I/O Server.
- ▶ Firmware update function of AIX diagnostics
- ▶ Firmware update function of the stand-alone *Diagnostics* CD
- ▶ Firmware update function of the *IBM Installation Toolkit for Linux* (IBMIT) DVD

This section shows in details all steps one needs to accomplish in order to update IBM BladeCenter JS23 and JS43 Express firmware image.

12.1.1 Commit new firmware to TEMP side

IBM BladeCenter JS23 and JS43 Express hold two firmware images in their flash, one located in the Temporary (TEMP) side, and the other located in the Permanent (PERM) side. Firmware updates *always* take place on the TEMP side.

Follow the steps below in order to commit a newer firmware image to your IBM BladeCenter JS23 and JS43 Express:

1. Make sure you are running the firmware locate in the TEMP side. See “Starting the firmware image from the TEMP side” on page 462 for more details.

Important: Before the installation of the new firmware to the TEMP side begins, the contents of the TEMP side should be committed to the PERM side.

The process provides for a *last known good copy of firmware* that can be used if the new image on the TEMP becomes corrupted.

2. Get the your new firmware image from IBM's web site. See “Get the latest firmware image from IBM” on page 471 for more details.

3. Copy the new firmware image file to your system, inside the `/tmp/fwupdate` or `/home/padmin/fw` for a VIO Server directory. You should create this directory if it doesn't exist. In order to do that type `mkdir /tmp/fwupdate` or `mkdir fw` for a VIO Server.
4. Log on to the AIX or Linux system as root, or log on to the Virtual I/O Server/IVM alpha partition as padmin.

Important: Updates from within an LPAR are not supported. You need to be logged in to the VIOS instead.

5. Type `ls /tmp/fwupdate` to confirm your new firmware image file is available in your system. This command will list all files inside the `/tmp/fwupdate` directory, including any firmware images you copied there. For this example we will use the name `01EAXXX.YYY.ZZZ.img` for our firmware image file. See "Interpreting the firmware file name" on page 470 for more details on how to interpret the numbers and letters of a real firmware file name.
6. Select one of the following methods to update the firmware image:
 - Install the firmware with the in-band diagnostics of AIX.
 - Update the firmware with the `update_flash` command on AIX

Example 12-1 Using the update_flash command inside AIX

```
cd /tmp/fwupdate  
/usr/lpp/diagnostics/bin/update_flash -f 01EAXXX.YYY.ZZZ.img
```

- Update the firmware with the `update_flash` command on Linux

Example 12-2 Using the update_flash command inside Linux

```
cd /tmp/fwupdate  
/usr/sbin/update_flash -f 01EAXXX.YYY.ZZZ.img
```

- Update the firmware with the `ldfware` command on VIOS

Example 12-3 Using the ldfware command inside VIOS

```
cd /tmp/fwupdate  
ldfware -file /home/padmin/fw/01EAXXX.YYY.ZZZ.img
```

7. Verify that the update was successful by checking the firmware level, as shown in "Verify current firmware level" on page 469.

12.1.2 Starting the firmware image from the TEMP side

Before running firmware updates, you need to make sure the BladeCenter server is using the firmware located in the TEMP side.

Note: Usually the IBM BladeCenter JS23 and JS43 Express are configured to use the TEMP side, leaving the firmware image in the PERM side as a backup.

It is possible to verify which side is being used, and change between firmware sides, from within the SMS menu, and the *Advanced Management Module* (AMM).

Configure to use the TEMP side through AMM

Perform the following steps in order to configure your BladeCenter server to use the firmware image on the TEMP side:

1. Access the Advanced Management Module.
2. Click **Blade Tasks** → **Configuration** → **Boot Mode**. Figure 12-1 on page 463 shows the AMM page for selecting blade's boot mode.

The screenshot shows the 'IBM BladeCenter H Advanced Management Module' interface. The left sidebar lists navigation options: Bay 1: xbch4amm, Monitors (System Status, Event Log, LEDs), Power Management, Hardware VPD, Firmware VPD, Remote Chassis, Blade Tasks (Power/Restart, Remote Control, Firmware Update Configuration, Serial Over LAN, Open Fabric Manager), I/O Module Tasks, MM Control, and Service Tools. The main content area is titled 'Boot Mode ?' and contains a table with 14 rows. The table columns are 'Bay', 'Name', 'Active Boot Mode', and 'Pending Boot Mode'. Rows 1 through 13 show blade servers with names like 'Free', 'E: Bad Console', 'JS23-Redbook', 'JS43-Redbook', and 'Japan_demo', all set to 'Temporary' in both modes. Row 14 shows 'No blade present'.

Bay	Name	Active Boot Mode	Pending Boot Mode
1	No blade present		
2	No blade present		
3	Free	Temporary	Temporary
4	No blade present		
5	No blade present		
6	E: Bad Console	Temporary	Temporary
7	No blade present		
8	JS23-Redbook	Temporary	Temporary
9	JS43-Redbook	Temporary	Temporary
10			
11	No blade present		
12	No blade present		
13	Japan_demo	Temporary	Temporary
14			

Figure 12-1 Select BladeCenter boot mode main page

3. Select the desired JS23 or JS43 blade server.
4. Select *Temporary* to force the system to use the firmware image from the TEMP side, as shown in Figure 12-2 on page 464, then click in **Save**.

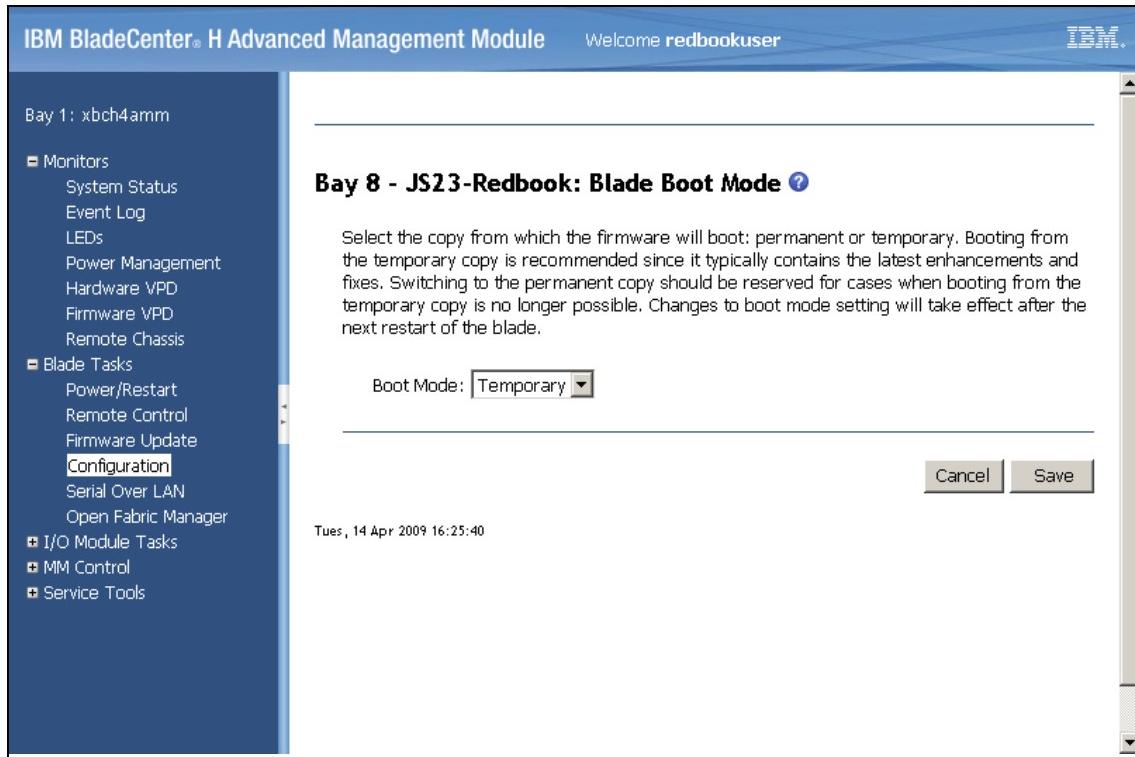


Figure 12-2 Firmware selection page

5. Restart the blade server. Click **Blade Tasks → Power/Restart**. Select the desired BladeCenter server in the list, then choose *Restart Blade* in the **Available Options** combobox. Finally, click **Perform Action**. Figure 12-3 on page 465 shows the Blade Power/Restart page.

IBM BladeCenter® H Advanced Management Module

Welcome redbookuser | About | Help | Logout

IBM

Bay 1: xbch4amm

- Monitors
 - System Status
 - Event Log
 - LEDs
 - Power Management
 - Hardware VPD
 - Firmware VPD
 - Remote Chassis
- Blade Tasks
 - Power/Restart**
 - Remote Control
 - Firmware Update
 - Configuration
 - Serial Over LAN
 - Open Fabric Manager
- I/O Module Tasks
- MM Control
- Service Tools

Blade Power / Restart ?

Blade selection and status

Click the checkboxes in the first column to select one or more blades; then, click one of the actions in the action list below the table and click Perform Action to perform the desired action.

	Bay	Name	Pwr	Local Pwr Control	Wake on LAN	Console Redirect
<input type="checkbox"/>	1	No blade present				
<input type="checkbox"/>	2	No blade present				
<input type="checkbox"/>	3	Free	Off	Enabled	N/A	
<input type="checkbox"/>	4	No blade present				
<input type="checkbox"/>	5	No blade present				
<input type="checkbox"/>	6	F: Bad Console	Off	Enabled	N/A	
<input checked="" type="checkbox"/>	7	No blade present				
<input checked="" type="checkbox"/>	8	JS23-Redbook	On	Enabled	N/A	
<input type="checkbox"/>	9	JS43-Redbook	On	Enabled	N/A	
<input type="checkbox"/>	10	No blade present				
<input type="checkbox"/>	11	No blade present				
<input type="checkbox"/>	12	No blade present				
<input type="checkbox"/>	13	Japan_demo	On	Enabled	N/A	
<input type="checkbox"/>	14					

Available actions

Figure 12-3 Blade Power / Restart

- Verify that the system starts using the firmware image from the TEMP side. It can be done by running steps 1 and 2 again (see Figure 12-1 on page 463).

Configure to use the TEMP side through the SMS menu

- Boot your blade server and hit 1 to enter the SMS menu, as shown in Figure 12-4 on page 466.

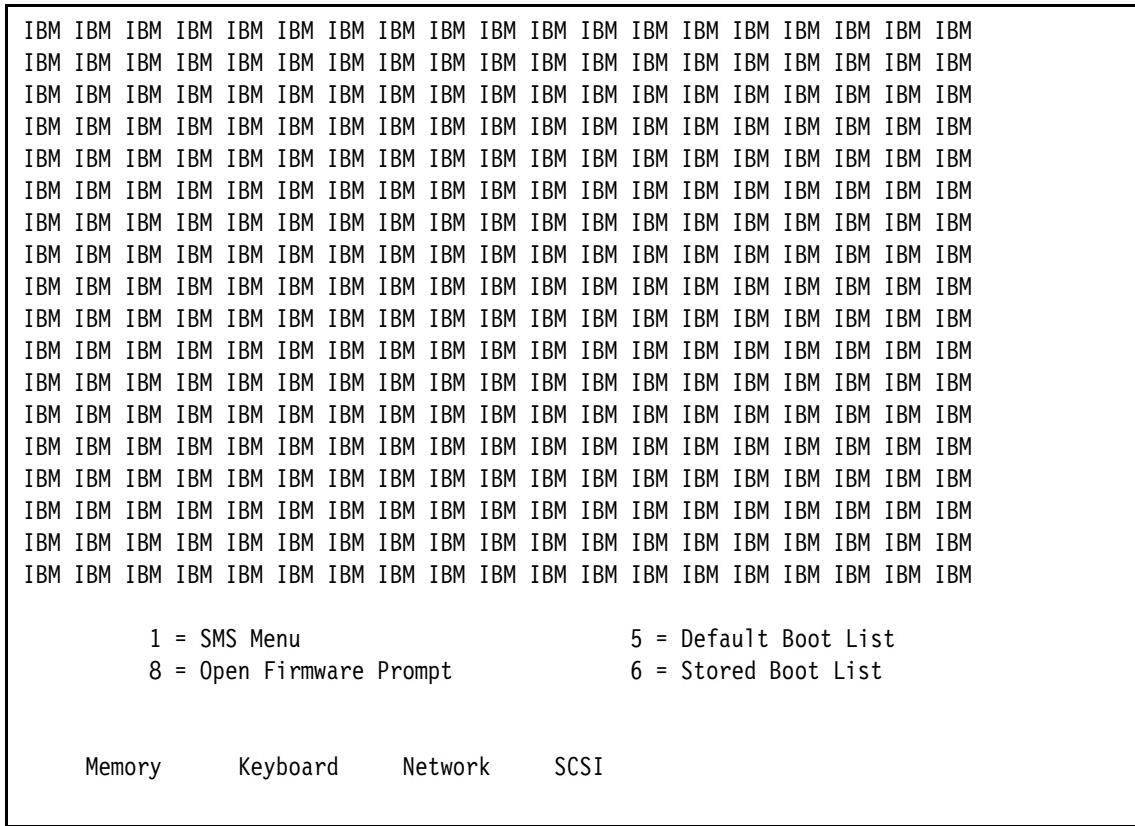


Figure 12-4 Enter SMS Menu

Note: Pay attention to the welcome screen shown in Figure 12-4 on page 466. It has a short time out, and if you miss it you'll need to reboot the machine.

2. Select **6 (Firmware Boot Side Options)**. As shown in Figure 12-5 on page 467.

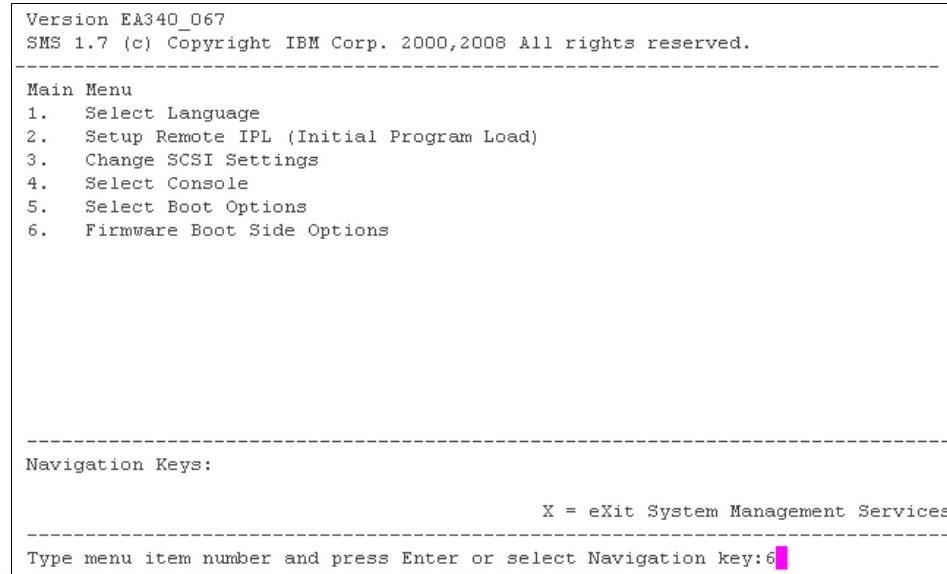


Figure 12-5 SMS main menu

Important: If your SMS menu does not provide option number 6, it is probably the case you are inside an LPAR. You can't run firmware updates in IBM BladeCenter JS23 and JS43 Express blade servers from within an LPAR.

3. Figure 12-6 on page 468 shows the SMS Boot Side Option Menu. In the upper left corner you can find the level of firmware being used, and just above options 1 and 2 you can find the firmware side being used. Type 2 and hit Enter to select the TEMP side for the next boot.

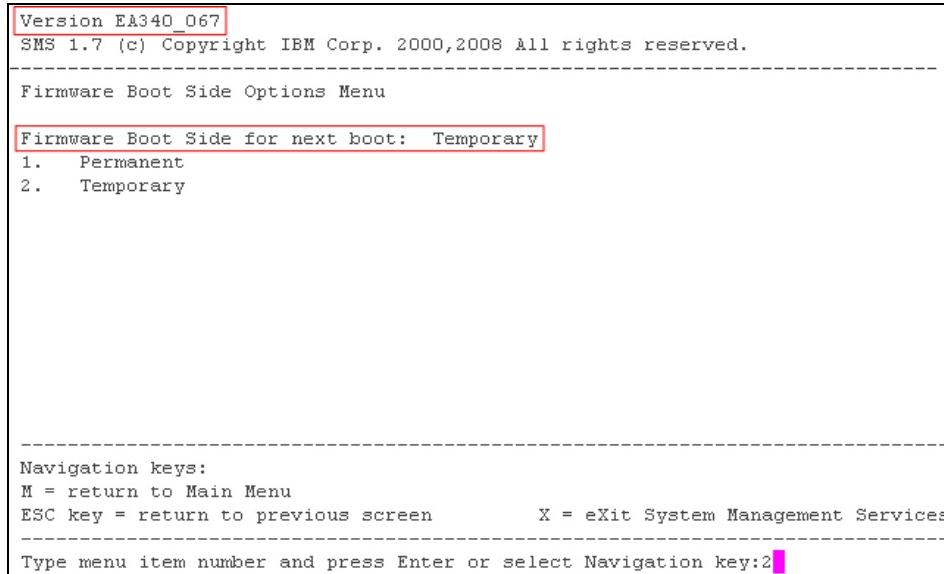


Figure 12-6 SMS firmware boot side options

4. Press X → 1 to restart the system, as shown in Figure 12-7.

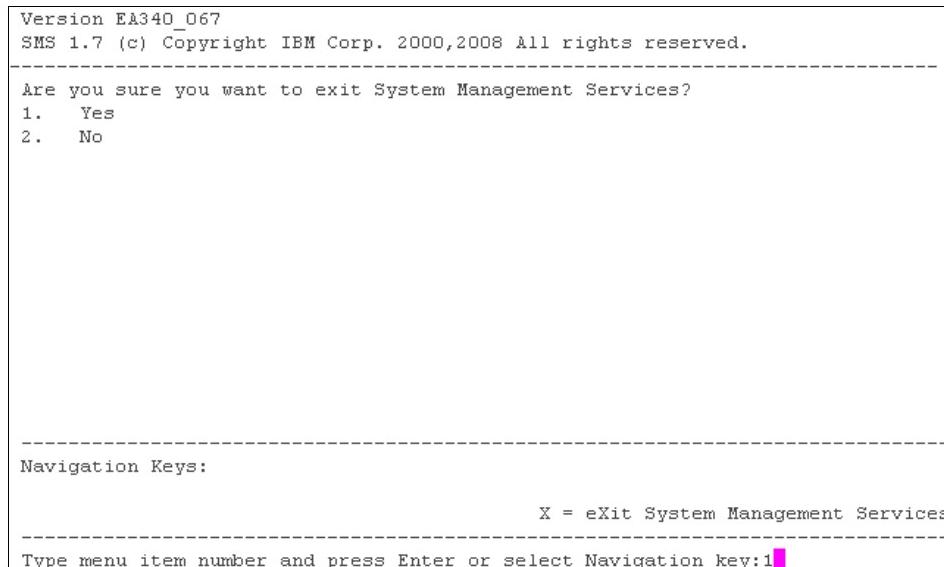


Figure 12-7 SMS exit page

12.1.3 Verify current firmware level

Before doing firmware updates, you must know which firmware level you are running in your IBM BladeCenter JS23 or JS43 Express. There are many ways to get this information, such as:

- ▶ Get firmware level through AMM.
- ▶ Get firmware level through SMS menu.
- ▶ Get firmware level through `lsmcode` command for Linux and AIX, or `lsfware` for Virtual I/O Server.

Get firmware level using the AMM

From within the AMM, click **Monitors** → **Firmware VPD**. Figure 12-8 shows the list with all blade servers and their firmwares.

The screenshot shows the IBM BladeCenter H Advanced Management Module (AMM) interface. The left sidebar has a tree view with nodes like 'Monitors' (selected), 'Blade Tasks', 'I/O Module Tasks', 'MM Control', and 'Service Tools'. The main content area has a title 'BladeCenter Firmware Vital Product Data' with a help icon. Below it is a link to 'Use the following links to jump down to different sections on this page.' followed by several hyperlinks: 'Blade Firmware Vital Product Data', 'I/O Module Firmware Vital Product Data', 'Management Module Firmware Vital Product Data', 'Power Module Cooling Device Firmware Vital Product Data', and 'Chassis Cooling Device Firmware Vital Product Data'. At the bottom of the content area, there's a note: 'To reread firmware Vital Product Data for a blade, select the blade, and click "Reload VPD". This process may take a while.' Below this is a 'Target' dropdown set to 'All Blades' and a 'Reload VPD' button. The central part of the screen displays a table titled 'Blade Firmware Vital Product Data' with columns: Bay (s), Name, Firmware Type, Build ID, Released, and Revision. The table lists the following data:

Bay (s)	Name	Firmware Type	Build ID	Released	Revision
3	Free	FW/BIOS	EA320_046	05/29/08	0818
6	F: Bad Console	Blade Sys Mgmt Processor	BOBT001		1.10
		Blade Sys Mgmt Processor	BOBT001	05/29/08	0818
8	JS23-Redbook	FW/BIOS		Unable to read blade VPD.	
9-10	JS43-Redbook	Blade Sys Mgmt Processor	EA340_067	03/17/09	0912
		Blade Sys Mgmt Processor	BOBT001		3.42
13-14	Japan_demo	FW/BIOS	EA340_068	03/19/09	0912
		Blade Sys Mgmt Processor	BOBT001		3.42

Figure 12-8 Firmware Vital Product Data (VPD) on AMM

Get firmware level using the SMS menu

1. Boot your blade server and hit 1 to enter the SMS menu, as shown in Figure 12-4 on page 466.

Note: Pay attention to the welcome screen shown in Figure 12-4 on page 466. It has a short time out, and if you miss it you'll need to reboot the machine.

2. Figure 12-9 shows the SMS main menu. In the left upper corner you can find the current firmware level.

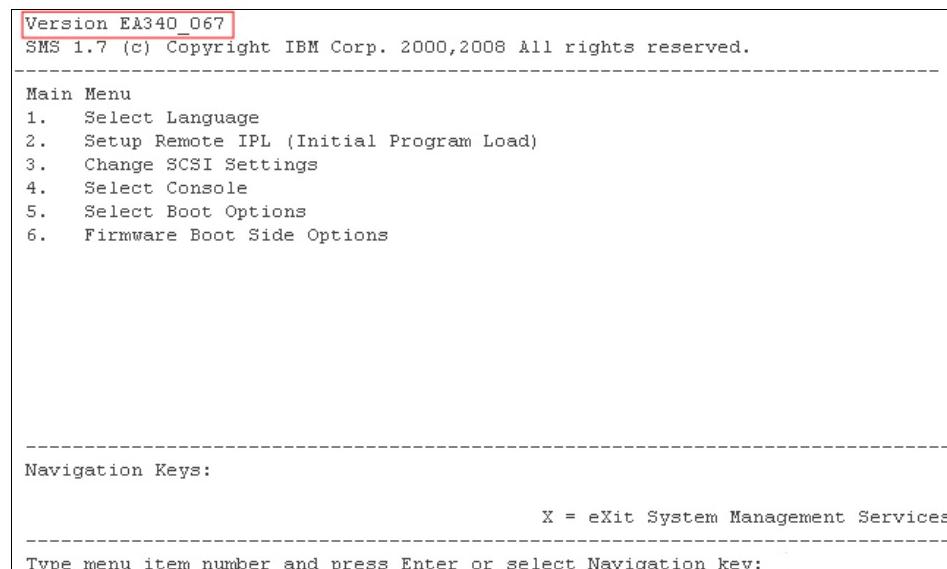


Figure 12-9 Firmware level inside the SMS main menu

12.1.4 Interpreting the firmware file name

Firmware image files for IBM BladeCenter JS23 and JS43 Express are the same of the previous JS12 and JS22 BladeCenter generation. They have the following versioning format:

EAXXX.YYY.ZZZ

where:

EA	It is an identifier. All IBM BladeCenter JS12, JS22, JS23 and JS43 Express firmware files will be have the EA prefix.
----	---

XXX	It is the <i>release level</i> . Changes in the release level means major updates in the firmware code.
YYY.ZZZ	They are the <i>service pack level</i> and <i>last disruptive service pack level</i> . Values for the service pack and last disruptive service pack are only unique within a release level.

A firmware installation is always disruptive if:

- ▶ New firmware release level is different from current firmware release level.
- ▶ New firmware service pack level and last disruptive service pack level have the same value.
- ▶ Current firmware's service pack level is lower than the last disruptive service pack level from the new firmware.

IBM BladeCenter server firmwares have a specific file name notation, which differs from other POWER systems. For a detailed description on the file name format, see:

ftp://ftp.software.ibm.com/systems/support/system_x_pdf/ibm_naming_convention.pdf

12.1.5 Get the latest firmware image from IBM

In order to do a firmware update in your IBM BladeCenter JS23 and JS43 Express, you need to get the firmware image file from IBM's BladeCenter Support web site
[\(http://www-947.ibm.com/systems/support/supportsite.wss/brandmain?brandind=5000020\)](http://www-947.ibm.com/systems/support/supportsite.wss/brandmain?brandind=5000020)

Important: In case you have problems accessing the link above, you can use the search function in IBM's main web page to find the correct URL.
Search for *Fix Central* or *Support and Downloads*.

Once inside the BladeCenter support site, select the appropriate Product family, Type, Model, and Operating System, to match your system's configuration. Click on **Go** when ready. Only the Product Family value is required.

You will be provided a list with all downloads available for your system. At this time we are looking for the BIOS files (which are actually the available firmware images for your system).

Click on the firmware version you want to download. You'll be redirected to a new page with many available file formats. For this example we will get the *Image File* format (with the *.img* extension).

12.2 System diagnostics

POWER6 processor-based systems contains specialized hardware detection circuits for detecting erroneous hardware operations, and includes extensive hardware and firmware recovery logic. IBM hardware error checkers have these distinct attributes:

- ▶ Continuous monitoring of system operations to detect potential calculation errors.
- ▶ Attempted isolation of physical faults based on runtime detection of each unique failure.
- ▶ initiation of a wide variety of recovery mechanisms designed to correct a problem.

Machine checks are handled by firmware. When a machine check occurs, the firmware analyses the error to identify the failing device and creates an error log entry.

In partitioned mode, any error that occurs during partition activity is surfaced to the operating system running in the partition. If some error occurs during POWER hypervisor (PHYP) activities, then the system gets rebooted by PHYP.

In case the system degraded to the point where the service processor cannot reach standby state, then the ability to analyze the error does not exist.

Important: This section is not intended to be a replacement for the information provided in the *BladeCenter JS23 and BladeCenter JS43 Type 7778 Problem Determination and Service Guide*, Part Number: 44R5339. For detailed steps on how to perform diagnostics tasks, determine the root cause of an error, and get proper support assistance, refer to this manual.

12.2.1 Diagnostic tools

This section brings a list of some tools that can be used to help in diagnostic hardware problems on IBM BladeCenter JS23 and JS43 Express.

Checkpoints and error codes

During system power-on process, the *Power-on self-test* (POST) checks out the hardware, including some system components and interconnections, and generates 8-digits checkpoint codes to mark the power-on progress.

If the Service Processor detects a problem during POST, an error code is logged in the AMM event log. Error codes are also logged in Linux syslog or AIX diagnostics log, if possible.

See “Checkpoint code (progress code)” on page 479 for more details.

Light Path and Front Panel diagnostics

IBM BladeCenter JS23 and JS43 Express comes with the Light Path technology, which helps on determining *Customer Replaceable Units* (CRU) with problems. Both the Base and MPE planar boards have LEDs for signaling faulty parts. The Front Panel is also part of the Light Path diagnostics solution, with LEDs to help verifying the current state of the system.

Front Panel (Control Panel)

IBM BladeCenter JS23 and JS43 Express have the same Front Panel, shown in Figure 12-10 on page 473.

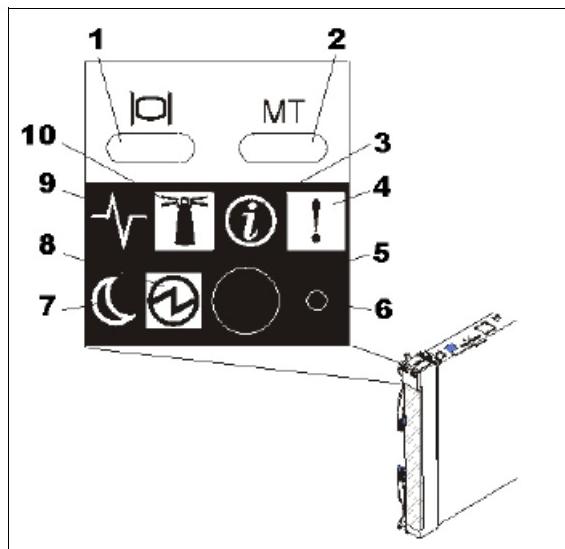


Figure 12-10 IBM BladeCenter JS23 and JS43 Control Panel

The Front Panel has buttons and LEDs that help on controlling and getting status from your Blade server. Table 12-1 has a short description for each button and LED located in the Front Panel. See *BladeCenter JS23 and BladeCenter JS43 Type 7778 Problem Determination and Service Guide*, Part Number: 44R5339 for more information.

Table 12-1 Description of Front Panel buttons and LEDs

Callout	Description
1	<i>Keyboard/Video</i> selection button.
2	<i>Media Tray</i> selection button.
3	<i>Information LED</i> .
4	<i>Error LED</i> .
5	<i>Power Control</i> button.
6	<i>Nonmaskable Interrupt (NMI) reset</i> button
7	<i>Sleep LED</i> . Not used in the IBM BladeCenter JS23 and JS43 Express.
8	<i>Power-on LED</i> .
9	<i>Activity LED</i> . When lit (green), it indicates that there is activity on the hard disk drive or network.
10	<i>Location LED</i> .

It is possible to control some of Front Panel LEDs through the management module, and through IBM Director Console. For instance, a system administrator can turn on or off the Location LED of some blade server.

The following steps show how to access these functions in the Advanced Management Module web interface:

1. Access the AMM web page.
2. Click **Monitors** → **LEDs**. Figure 12-11 on page 475 shows the LEDs monitoring page in AMM.

Bay	Name	Pwr*	Error	Information	KVM	MT	Location
1	No blade present	Off	Off	Off	Off	Off	On Off Blink
2	No blade present	Off	Off	Off	Off	Off	On Off Blink
3	Free	Off	Off	Off	Off	Off	On Off Blink
4	No blade present	Off	Off	Off	Off	Off	On Off Blink
5	No blade present	Off	Off	Off	Off	Off	On Off Blink
6	F: Bad Console	Off	Off	Off	Off	Off	On Off Blink
7	No blade present	Off	Off	Off	Off	Off	On Off Blink
8	JS23-Redbook	On	!	Off	Off	On	Off Blink
9	JS43-Redbook	On	!	Off	Off	On	Off Blink
11	No blade present	Off	Off	Off	Off	Off	On Off Blink
12	No blade present	Off	Off	Off	Off	Off	On Off Blink
13	Japan_demo	On	!	Off	Off	Off	On Off Blink
14	Japan_demo	On	!	Off	Off	Off	On Off Blink

* If a blade is powered off, its physical LEDs are not lit. This table represents the status of all LEDs, even for powered-off blades.

Figure 12-11 AMM BladeCenter LEDs control and status page

Light Path

Light Path diagnostics is a system of LEDs on the control panel and on your system board (IBM BladeCenter JS43 Express has Light Path LEDs on both boards). When a hardware error occurs, LEDs are lit throughout the blade server.

LEDs are available for many components, such as:

- ▶ Battery.
- ▶ SAS HDD (or SSD) disks, on both Base and MPE planars.
- ▶ Management card on Base planar only.
- ▶ Memory modules on both Base and MPE planars.
- ▶ Expansion cards (1Xe and HSDC daughter cards) on both Base and MPE planars.

The MPE planar has a special LED to help on determining the correct planar where the error is happening. This LED is called “*Check card below*”, and is used to indicate that the problem is not occurring in the MPE planar, but in the Base planar. This LED is also called “*Not me*” LED.

Note: We recommend you to see the *BladeCenter JS23 and BladeCenter JS43 Type 7778 Problem Determination and Service Guide*, Part Number: 44R5339. There you will find more detailed information on how to perform diagnostics using the Light Path technology, and also how to act when some well-known types of problems arise.

Figure 12-12, Figure 12-13 on page 477, and Table 12-2 on page 477 show all Light Path LEDs available on your IBM BladeCenter JS23 and JS43 Express boards.

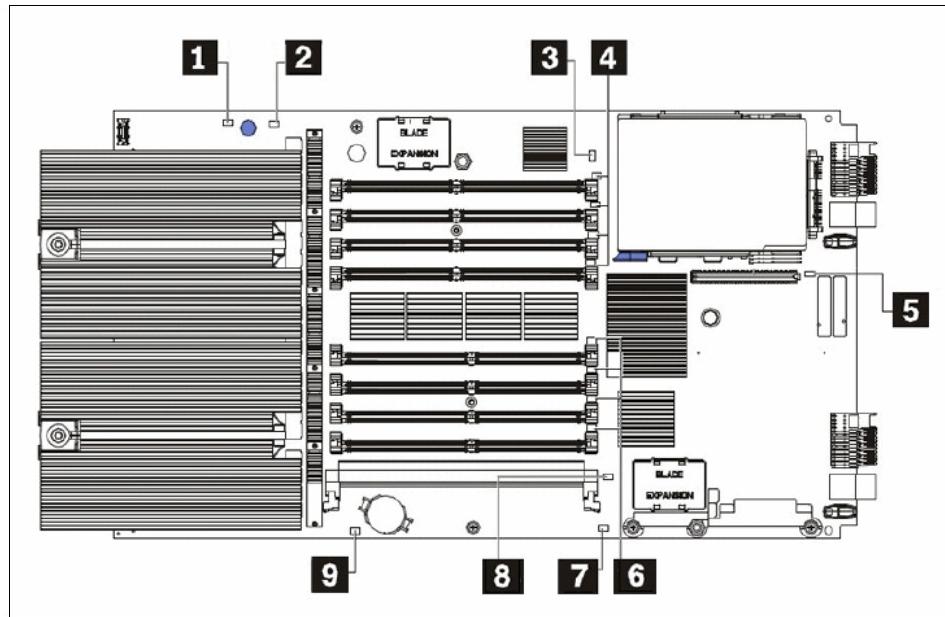


Figure 12-12 LEDs on the IBM BladeCenter JS23 Express board. This is the same LED configuration of the IBM BladeCenter JS43 Express Base planar

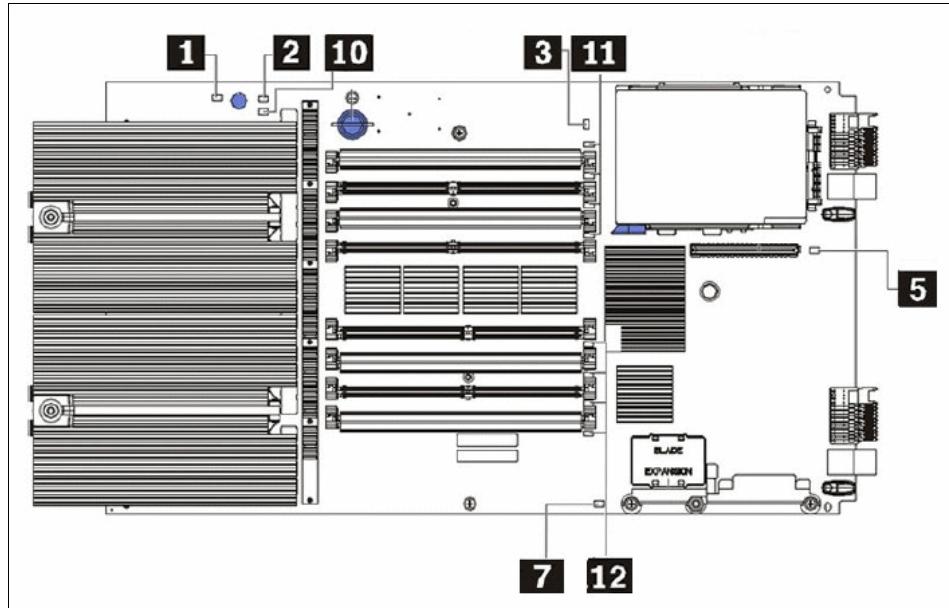


Figure 12-13 LEDs on the IBM BladeCenter JS43 Express MPE planar

Table 12-2 Lightpath LED description.

Callout	Base planar LEDs	Expansion planar LEDs
1	Lightpath power LED	
2	System board LED	
3	SAS hard disk drive LED	
4	DIMM 1-4 error LEDs	None
5	1Xe expansion card connector LED	
6	DIMM 5-8 error LEDs	None
7	PCIe high-speed expansion card LED	
8	Management Card LED	None
9	Battery LED	None
10	None	<i>“Check card below” LED</i>
11	None	DIMM 9-12 error LEDs
12	None	DIMM 13-16 error LEDs

Diagnostic utilities for the AIX operating system

AIX provides many diagnostic and maintenance functions, such as:

- ▶ Automatic error log analysis.
- ▶ Firmware updates, format disk, and RAID Manager.

For more information on how to perform diagnostics in your IBM BladeCenter JS23 and JS43 Express using AIX, please see

<http://publib.boulder.ibm.com/infocenter/systems/scope/hw/index.jsp?topic=/iphau/working01.htm>

Diagnostic utilities for the Linux operating system

Linux on POWER service and productivity tools (RAS Tools) include hardware diagnostic aids and productivity tools for Linux on POWER systems.

In order to do error log analysis in Linux, you need to have the *diagela* package installed.

Note: SuSE Linux Enterprise Server 11 (SLES11) ships a new package called *ppc64-diag* for running log analysis on POWER machines. You should install *ppc64-diag* package instead of *diagela* on SLES11.

See Appendix , “IBM service diagnostic aids and productivity tools overview” on page 546 for more details on the RAS Tools for Linux.

Diagnostics utilities for the IBM i operating system

IBM i only runs in logical partitions on IBM BladeCenter JS23 and JS43 Express, therefore it sees virtual devices only. In this situation, the errors associated with hardware failures will not be seen inside the operating system, in the normal Product Activity Logs (PAL) or System Activity Logs (SAL).

See “IBM i partition diagnostics and errors” on page 483 for a detailed description on how to gather diagnostic information from an IBM i partition.

Stand-alone diagnostics

When you don't have any operating system installed in your Blade server, or you are not able to boot the installed operating systems, it is still possible to run diagnostic utilities using one of these Stand-alone diagnostic solutions:

- ▶ AIX-based Stand-alone *Diagnostics* CD, which includes all diagnostic utilities found in AIX.
- ▶ *IBM Installation Toolkit for Linux* DVD, which can be used as a Linux Live DVD, and includes all RAS Tools for Linux. See Appendix , “IBM service

diagnostic aids and productivity tools overview” on page 546 for more details on the IBM Installation Toolkit for Linux.

12.2.2 Reference codes

Reference codes are diagnostic aids that help you determine the source of a hardware or operating system problem. IBM BladeCenter JS23 and JS43 Express produces many types of Reference Codes:

- | | |
|----------------|---|
| Progress codes | 8-digit status codes generated by the Power-on Self-test (POST). They are used to show progress when powering on the blade server. |
| Error codes | 9-word (8-digit per word) error code. Error codes are either <i>System Reference Codes (SRCs)</i> or <i>Service Request Numbers (SRN)</i> . It can also have a <i>Location Code</i> , which helps on isolating the piece of hardware with problems. |

The following subsections give a brief description on the various types of codes generated by the IBM BladeCenter JS23 and JS43 Express. For a more detailed description of each code type, and a list of possible codes, see the *BladeCenter JS23 and BladeCenter JS43 Type 7778 Problem Determination and Service Guide*, Part Number: 44R5339.

For details on how to visualize the Reference codes for a specific blade server, see: “Diagnostics tasks through the Advanced Management Module” on page 480.

Location code

Location codes identify components of the blade server. They can appear in some error codes to help identifying which server’s component is causing the error.

There are Location codes for components in both Base and MPE planars.

Checkpoint code (progress code)

Checkpoint codes are generated by the power-on self-test (POST), which performs a series of tests to check the operation of the blade server components.

Progress codes do not indicate an error, although in some cases, the blade server can pause indefinitely (hang). Progress codes for blade servers are 9-word, 8-digit hexadecimal numbers that start with C and D.

A checkpoint might have an associated location code as part of the message. The location code provides information that identifies the failing component when there is a hang condition.

System reference code (SRC)

System reference codes are used to identify both hardware and software problems in IBM BladeCenter JS23 and JS43 Express. Those errors can be originated in hardware, in firmware, or in the operating system.

The SRC identifies the component that generated the error code and describes the error. Each SRC has nine words, each word being an eight-digit code.

Usually the first character indicates the type of error, as shown in Table 12-3

Table 12-3 Types of SRCs

First character	Type of error
A	User Attention, which means that user interventions may required to progress.
B	Error. Firmware has identified an error.
C	Checkpoint. Used to give status on normal boot. Hang indicates unexpected error.
D	Dump. Indicates a dump is in progress. Hang indicates unexpected error.
1	<i>System Power Control Network (SPCN) error.</i>
6	Virtual optical device error.

Service request number (SRN)

Service request numbers are error codes generated by the operating system, and they have three digits, a hyphen, then three or four digits after the hyphen. This type of error code can be visualized by the AIX *diagnostics* or the Linux service aid *diagela* (*ppc64-diag* on SLES11) package if it is installed.

12.2.3 Diagnostics tasks through the Advanced Management Module

This section describes how to use some AMM features to diagnostic issues in your IBM BladeCenter JS23 and JS43 Express. For detailed information on how to install, configure, and use the AMM, see:

- ▶ Advanced Management Module User's Guide
<ftp://ftp.software.ibm.com/systems/support/intellistation/44r5375.pdf>
- ▶ Advanced Management Module Installation Guide
ftp://ftp.software.ibm.com/systems/support/system_x/44r5269.pdf

Using the AMM to view SRCs

You can use the AMM to view the progress codes and error codes for a specific blade server, just follow the steps below:

Note: IBM BladeCenter JS23 and JS43 doesn't display the checkpoints in the Serial-Over-Lan (SOL) console, neither in the graphical console.

1. Access the AMM web interface.
2. In the main menu on the left, click **Service Tools → Blade Service Data**, as shown in Figure 12-14 on page 481.

Bay	Name
1	No blade present
2	No blade present
3	Free
4	No blade present
5	No blade present
6	F: Bad Console
7	No blade present
8	JS23-Redbook
9	JS43-Redbook
10	
11	No blade present
12	No blade present
13	
14	Japan_demo

Figure 12-14 AMM Blade Server Data panel

3. Select the desired blade server. The Reference codes will be shown for the chosen blade server, as in Figure 12-15 on page 482.

The Advanced Management Module can display the last 32 Reference Codes. You can manually refresh the list to update it.

Unique ID	System Reference Code	Timestamp
000000ff	AA00E1A9	2009-04-17 19:11:54
000000fe	CA00E1A0	2009-04-17 19:11:54
000000fd	CA00E1B5	2009-04-17 19:11:54
000000fc	CA00E1F1	2009-04-17 19:11:50
000000fb	CA00E1F0	2009-04-17 19:11:50
000000fa	CA00E141	2009-04-17 19:11:50
000000f9	CA00E1DC	2009-04-17 19:11:42
000000f8	CA00E1DC	2009-04-17 19:11:42
000000f7	CA00D008	2009-04-17 19:11:42
000000f6	CA00E100	2009-04-17 19:11:42
000000f5	CA00E1FB	2009-04-17 19:11:42
000000f4	CA00E100	2009-04-17 19:11:42
000000f3	CA00E1FF	2009-04-17 19:11:42
000000f2	CA00E1FE	2009-04-17 19:11:42
000000f1	CA00E1FF	2009-04-17 19:11:42
000000f0	CA00E1FE	2009-04-17 19:11:42

Figure 12-15 Power-on checkpoints inside AMM web interface

Using the AMM to view log messages

You can use the AMM web interface to view log messages generated by the blade servers within a Blade Center chassis.

Once inside AMM, click **Monitors** → **Event Log**. Figure 12-16 on page 483 shows the event log interface in AMM.

It is possible to filter the log events by blade server, by severity of event, by date, and also by serviceable flag, that indicates whether the *Call Home* flag was enabled for that event or not. For more details on *Call Home*, see “Service Advisor” on page 483.

Index	Sev	Source	Date/Time	Text
1	I	Audit	04/21/09, 15:26:36	Remote login successful for user 'redbookuser' from Web at IP 9.41.30.158
2	I	Audit	04/21/09, 15:18:46	Remote login successful for user 'redbookuser' from TCP Cmd at IP 9.41.30.108

Figure 12-16 AMM event log interface

Service Advisor

The Service Advisor enables the Blade Center to automatically send hardware and firmware serviceability messages to IBM. Every time a log event with the Call Home flag enabled happens, the AMM's Service Advisor will send a message with the event log message, BladeCenter unit inventory and status to IBM Support.

This Call Home feature comes disabled by default. You need to accept the Service Advisor Terms and Conditions before enabling it.

For a detailed description on how to enable the Service Advisor in the AMM web interface, see

<ftp://ftp.software.ibm.com/systems/support/intellistation/44r5375.pdf>

Note: No client data from the server or I/O modules is sent within the Call Home contact message.

12.2.4 IBM i partition diagnostics and errors

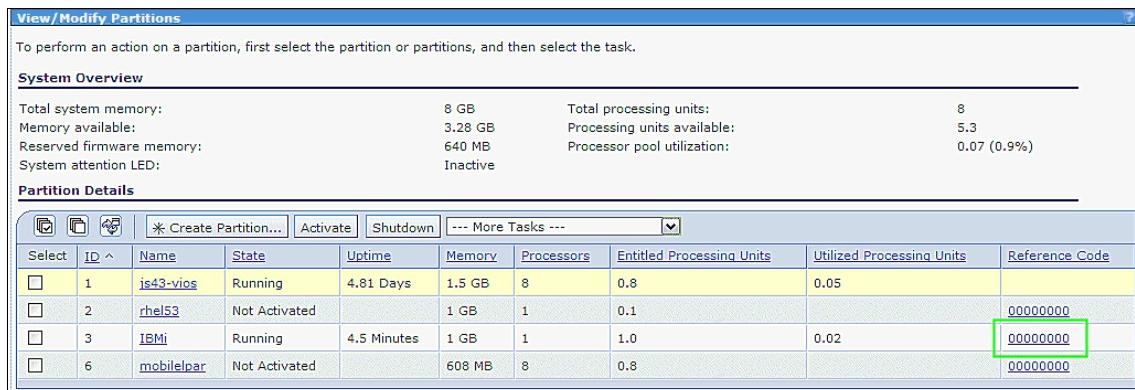
When an IBM i partition is configured on a JS23/JS43 it will be configured with virtual devices only. Since an IBM i partition on a blade server cannot own any physical resources the errors associated with hardware failures will not be seen

in the normal Product Activity Logs (PAL) or System Activity Logs (SAL). Most errors the IBM i partition will encounter are going to be related to storage or configuration. Any true hardware errors will be reported to the VIOS partition and repaired using VIOS options. In this section we will outline where to collect error data and configuration information related to an IBM i virtual partition.

Collecting IBM i partition error information

When an IBM i partition encounters a failure, the reference code information is what will be needed to troubleshoot and repair errors. The reference codes can be found using IVM on the View/Modify Partitions panel.

Figure 12-17 on page 484 shows an example of the partitions view with reference codes. The reference codes are in the far right column by default. Clicking on the reference code for the partition selected.



Select	ID	Name	State	Uptime	Memory	Processors	Entitled Processing Units	Utilized Processing Units	Reference Code
<input type="checkbox"/>	1	js43-vios	Running	4.81 Days	1.5 GB	8	0.8	0.05	
<input type="checkbox"/>	2	rhel53	Not Activated		1 GB	1	0.1		00000000
<input type="checkbox"/>	3	IBMi	Running	4.5 Minutes	1 GB	1	1.0	0.02	00000000
<input type="checkbox"/>	6	mobilepar	Not Activated		608 MB	8	0.8		00000000

Figure 12-17 Partitions view

Another entry point to the reference codes is to **select the partition** with the check box, then use the **More Tasks** pull down and select **Reference Codes**. Figure 12-18 on page 485 provides an example of this option.

System Overview

Total system memory:	8 GB	Total processing units:
Memory available:	3.28 GB	Processing units availab
Reserved firmware memory:	640 MB	Processor pool utilization
System attention LED:	Inactive	

Partition Details

Select	ID ^	Name	State	Uptime
<input type="checkbox"/>	1	js43-vios	Running	4.81 Days
<input type="checkbox"/>	2	rhel53	Not Activated	
<input checked="" type="checkbox"/>	3	IBMi	Running	7.4 Minutes
<input type="checkbox"/>	6	mobileipar	Not Activated	

More Tasks ---

- More Tasks ---
- Open terminal window
- Delete
- Create based on
- Operator panel service functions
- Reference Codes
-
- Mobility ---
- Migrate
- Status
-
- Properties

Figure 12-18 More Tasks - Reference Codes

Once the reference codes option is selected, a new window will appear that displays the list of codes for the partition selected. In Figure 12-19 on page 486 the codes listed are from the last IPL. Everything is normal with no errors at this time. Selecting any reference code will display the additional words to the right of the panel in the details section.

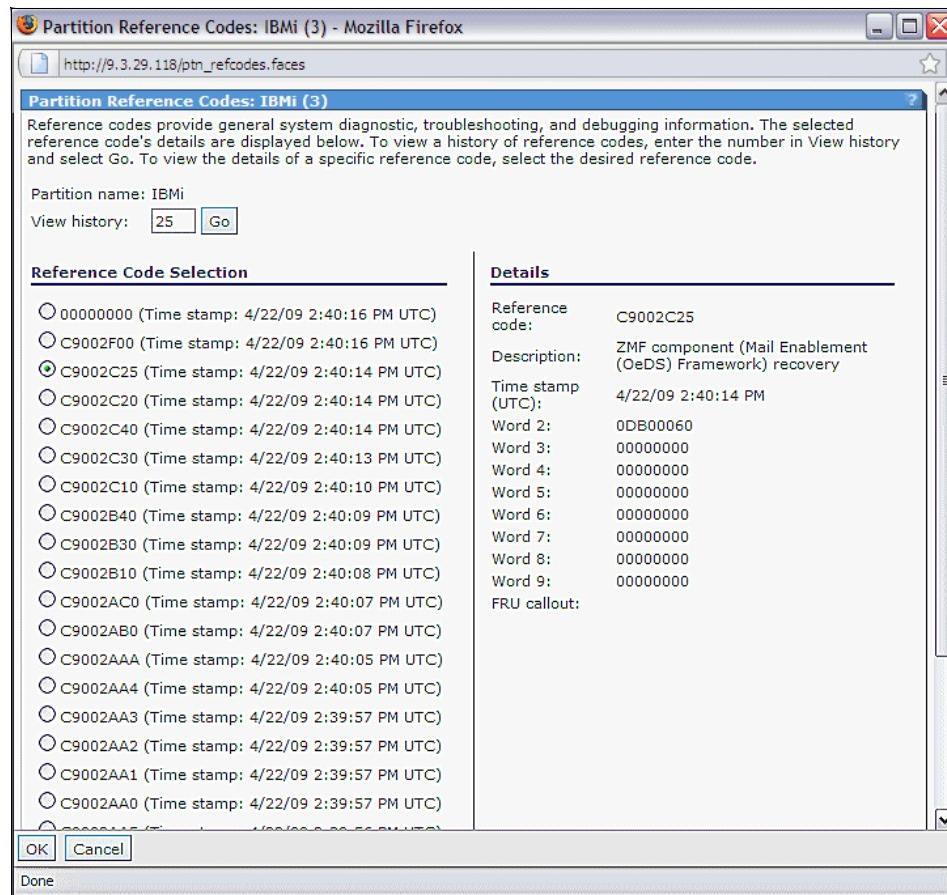


Figure 12-19 Reference Code list - normal IPL

Now let's look at an error condition in the IBM i partition. For this scenario we will assume the partition was running with no problems. Something happened that caused the partition to hang. Users report that the partition is not responding.

There are many ways to troubleshoot and report problems. It is not the intent of this section to provide procedures beyond collecting data and contacting your next level of support.

Depending on your systems configuration, the chassis can be configured to notify the user of errors. This type of error can be included in the notification type so the user may get an E-mail indicating an error on the Blade Center has occurred.

One of the places to look for errors will be in IVM. Looking at the View/Modify partitions screen we notice an error condition on the IBM i partition. In Figure 12-20 on page 487 notice that the Attention Indicator is next to the partition and in the reference code column there is a code listed. Normally we expect to see 00000000 in the reference code column if everything is running ok.

The screenshot shows the 'View/Modify Partitions' interface. The 'System Overview' section displays system memory (8 GB total, 3.28 GB available), processing units (8 total, 5.3 available), and a system attention LED that is active with a warning icon. The 'Partition Details' section lists six partitions:

Select	ID	Name	State	Uptime	Memory	Processors	Entitled Processing Units	Utilized Processing Units	Reference Code
<input type="checkbox"/>	1	is43-vios	Running	4.82 Days	1.5 GB	8	0.8	0.17	
<input type="checkbox"/>	2	rhel53	Not Activated		1 GB	1	0.1		00000000
<input type="checkbox"/>	3	IBMi	Running Attention	18.7 Minutes	1 GB	1	1.0	0.00	A6010266
<input type="checkbox"/>	6	mobilepar	Not Activated		608 MB	8	0.8		00000000

Figure 12-20 IBM i partition error

To determine more about the error click on the reference code or use the More Tasks option as mentioned earlier. Figure 12-21 on page 488 shows an example of the reference codes with the error listed.

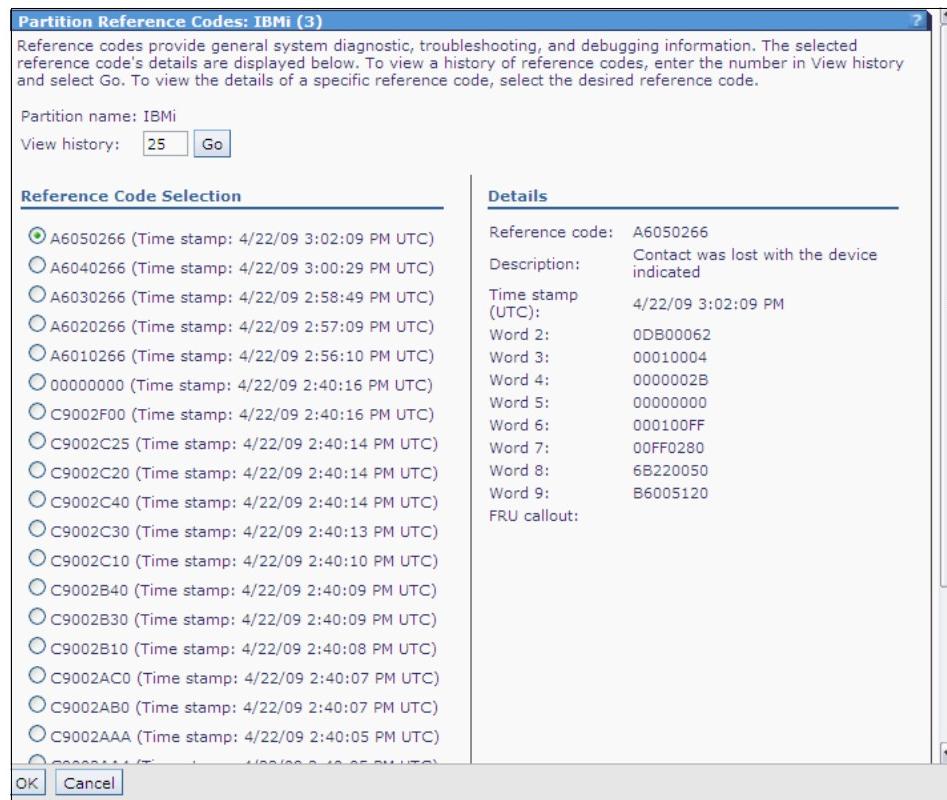


Figure 12-21 Reference Code list - error condition

Using the start of call procedures this reference code information would be used to complete the Problem Summary Form. This information would be used by service and support to troubleshoot the error and provide assistance in resolving the problem. Depending on your skill level, you may be able to navigate through the various Information Center pages to troubleshoot this error further.

Another source for error information would be from the AMM. As mentioned above, the AMM can be configured to send alerts on certain error conditions. Even if the AMM is not configured in this manner you can still use the Event Logs to view and collect error information.

Figure 12-22 on page 489 shows an example of the event log on the AMM.

Figure 12-22 AMM Event Log

The event log can be filtered to view only events specific to the blade server or other components. Figure 12-23 on page 489 shows an example of the filter options.

Figure 12-23 Event log filter

In the list of events you will see the error log information. Figure 12-24 on page 490 provides an example of the data in the AMM event log. This data should be similar to the data shown on the partition reference code screen as viewed

from IVM we looked at earlier. This data can also be saved by scrolling to the bottom of the event log and using the **Save Log as Text File** button. This data could then be supplied to service and support for further assistance in error determination.

x	Sev	Source	Date/Time	Event ID	Text
	E	Blade_09	04/22/09, 10:56:27	0x10000002	(JS43-Redbook) SYS F/W: Error. Replace SVCDOCS (5303B448 A6000266 00000062 00010004 00000028 00000000 000100FF 00FF0280 6B220050 B6005120)
	E	Blade_09	04/22/09, 10:54:48	0x10000002	(JS43-Redbook) SYS F/W: Error. Replace SVCDOCS (5303B446 A60F0266 00000062 00010004 00000028 00000000 000100FF 00FF0280 6B220050 B6005120)

Figure 12-24 Event log data details

As mentioned above, it is not the intention of this book to explain troubleshooting processes for an IBM i partition. For those that are curious about this particular error, it was created by removing the partition assignment of one of the hdisk units for the partition.

Typically errors of this type in a pure virtual environment are caused by configuration changes made by the administrator or others that might not be familiar with the partition configuration. In our example we just added the drive back into the configuration using the Modify Partition Assignment options and the partition recovered. Using IVM you can see the tasks performed by the Monitor Tasks option. This may provide some indication a configuration change was made that could have caused the error on the partition. Figure 12-25 on page 490 shows an example of the Monitor Task option.

Monitor Tasks					
The last 40 tasks that have been run from the web interface on the system are listed below. Select the task to view its properties.					
Monitored Tasks					
Select	Name	Status	Start Time	Task Duration	User ID
<input checked="" type="radio"/>	Modify_partition_assignment (physical volume)	Successful	4/22/09 11:04:13 AM	00:00:00	padmin
<input type="radio"/>	Modify_partition_assignment (physical volume)	Successful	4/22/09 9:55:53 AM	00:00:00	padmin
<input type="radio"/>	Activate	Successful	4/22/09 9:36:43 AM	00:00:00	padmin

Figure 12-25 Monitor tasks

While the task may not specifically indicate what exactly was modified, at least there is an indication something changed possibly around the time of the partition error.

Removing and adding drives to IBM i partitions can be disastrous to the partition and may cause system reload conditions if care is not exercised to follow the correct processes.

Appendices

In this part of the book we provide additional technical support information:

- ▶ Appendix A, “Consoles, SMS, and Open Firmware” on page 493
- ▶ Appendix B, “SUSE Linux Enterprise Server AutoYaST” on page 521
- ▶ Appendix C, “Additional Linux installation configuration options” on page 535
- ▶ Appendix D, “Service and productivity tools for Linux” on page 545



A

Consoles, SMS, and Open Firmware

This appendix briefly covers the methods to gain access to the console, use the System Maintenance Services Menu (SMS) to select the console to use, and use the Open Firmware prompt to choose fiber channel host bus adapter settings.

This Appendix has the following sections:

- ▶ “Consoles of the IBM BladeCenter JS23 and JS43” on page 494
- ▶ “System Management Services menu” on page 501
- ▶ “Open Firmware interface” on page 509

Consoles of the IBM BladeCenter JS23 and JS43

Like the previous JS12 and JS22 BladeCenter servers, the IBM BladeCenter JS23 and JS43 blades have a graphics adapter. This graphics adapter makes it possible to use the KVM switch that is built into the Advanced Management Module to gain access to the console of the blade.

An alternative method to gain access to the console is the usage of Serial Over LAN, or SOL.

You can use either the graphical console or the SOL console during POST. A JS23/JS43 blade will start the first time as default with the graphical console as long as you do not change this. The blade will ask, after first power on, which console should be used as active console. If you miss this point during POST you need to enter the System Maintenance Services menu from a graphical console and change the console. See “System Management Services menu” on page 501 for information about how to change the active console.

Graphical console

The graphical console is available by either using the connected keyboard, display, and mouse on a management module installed in the BaldeCenter chassis, or the remote control function of the Advanced Management Module.

Using keyboard, video, and mouse connected to a management module to gain access to the graphical console

The physical console connected to a management module will most of the time not be used. If you are working locally at the BladeCenter, you may use this console to access blades. There is one analog console available that is switched via the KVM switch that is built into the management module between the blades in the chassis. To switch the console you may use the keyboard/video select button that is placed on the front panel of each blade. Figure A-1 on page 495 shows the location of the front panel itself as well as the control elements on the panel of an IBM BladeCenter JS23 and JS43.

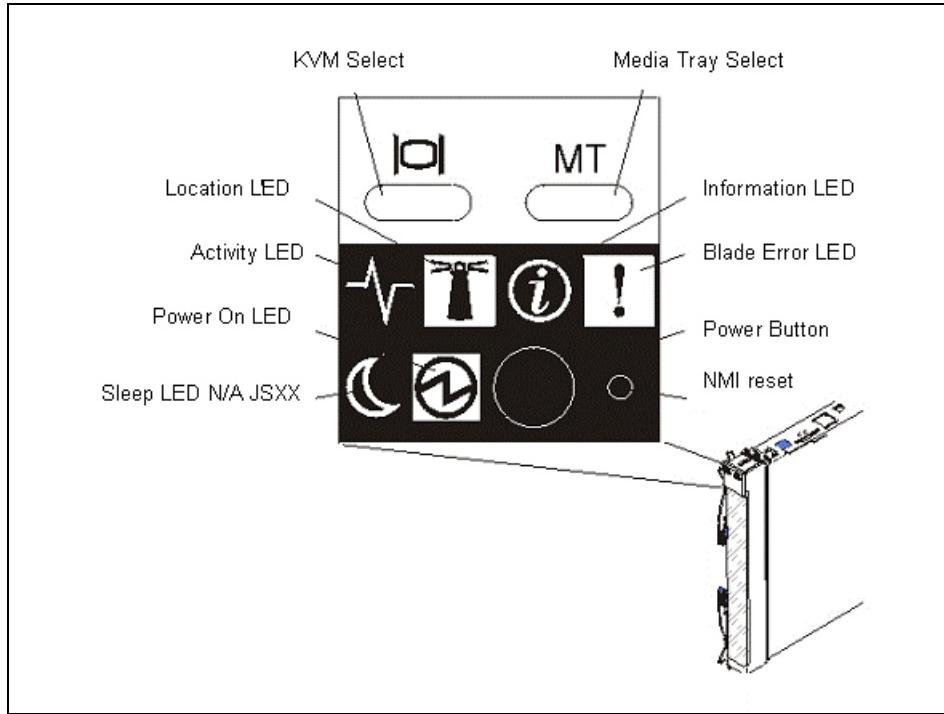


Figure A-1 JS23/JS43 Control Panel

Pressing the keyboard/video select button switches the console to the blade on which the button was pressed. There is only one blade in a chassis that has lit the keyboard/video select button.

Note: Be sure that you are using the keyboard, video, and mouse connected to the active Advanced Management Module. There is only one management module active at one time. You will recognize this by looking at the management modules' LEDs. Only one management module has the blinking heartbeat LED switched on. This applies to installations with redundant management modules only.

You may use the `list -1 2` command in a telnet or SSH session connected to the Advanced Management Module to identify the active management module.

An alternative method to switch the physical console between the blades is the usage of the key combination Shift+Num lock+Num lock+Bay Number+Enter on an IBM Space Saver Keyboard that is typically used in an IBM Rack. On standard keyboards the shift key is not required.

Use the key combination as follows:

1. Click and hold the Shift key.
2. Click Num Lock twice.
3. Release the Shift key.
4. Click the bay number - one of 1-14 depending on the chassis you are using.
5. Click Enter.

Using remote control to access the graphical console

Remote control is a feature of the management module installed in a BladeCenter chassis. It allows to connect over an IP connection to the management module and open a browser window that has the graphical console redirected. Only one analog console redirection can be in progress at one time. This remote control shows the same content that is shown on the physical console.

Note: Remote control supports at maximum the standard VESA modes with 1024x768 pixels at 60,70 and 75Hz.

Concurrent KVM (cKVM) allows multiple digital KVM connections on some blades. This feature is not supported on IBM BladeCenter JS23 and JS43.

1. To open a remote control session, log on to the management module with your Web browser. The default IP address of the management module is 192.168.70.125/24. Enter the user name and password and click **Log In** as shown in Figure A-2 on page 497. The default account with all administrative rights is USERID with password PASSW0RD.

Note: It is recommended that each administrator has its own account. Authentication against an LDAP server can be used when more accounts are required than are possible to configure on the management module.

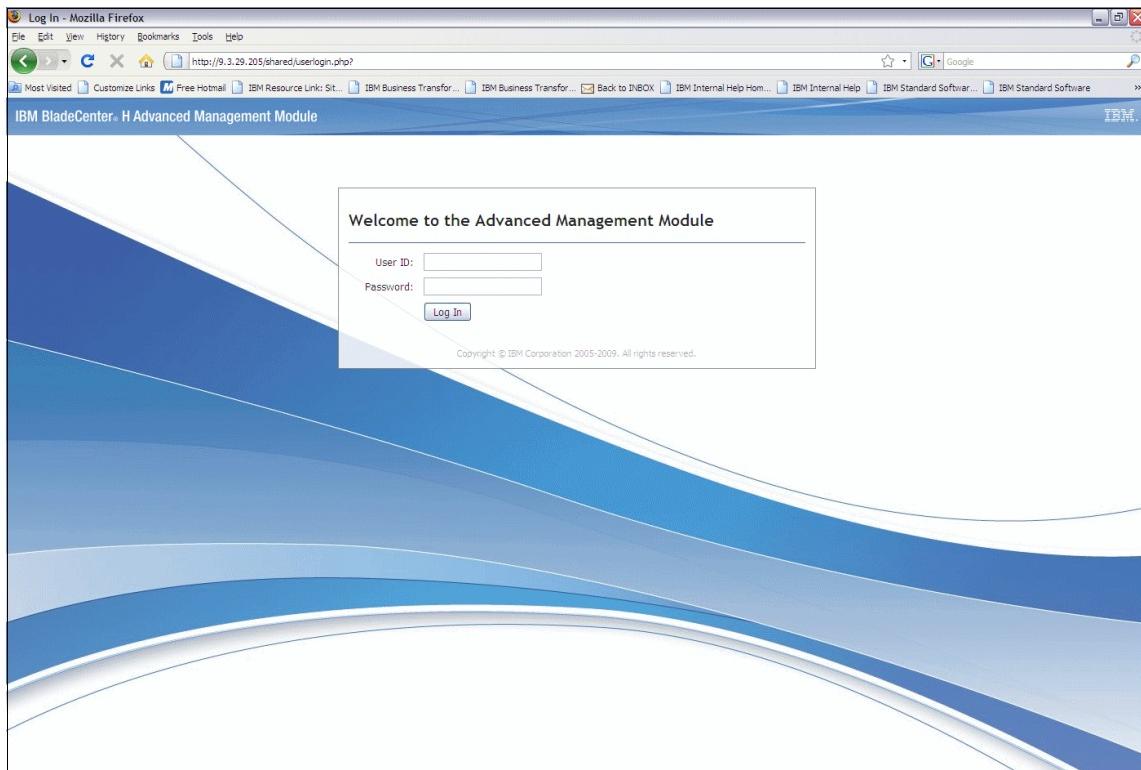


Figure A-2 AMM login panel

2. If prompted, select the time-out parameter that defines after how much idle time the session will be closed. Click **Continue**; Figure A-3. Our example has been modified to show no time-out.

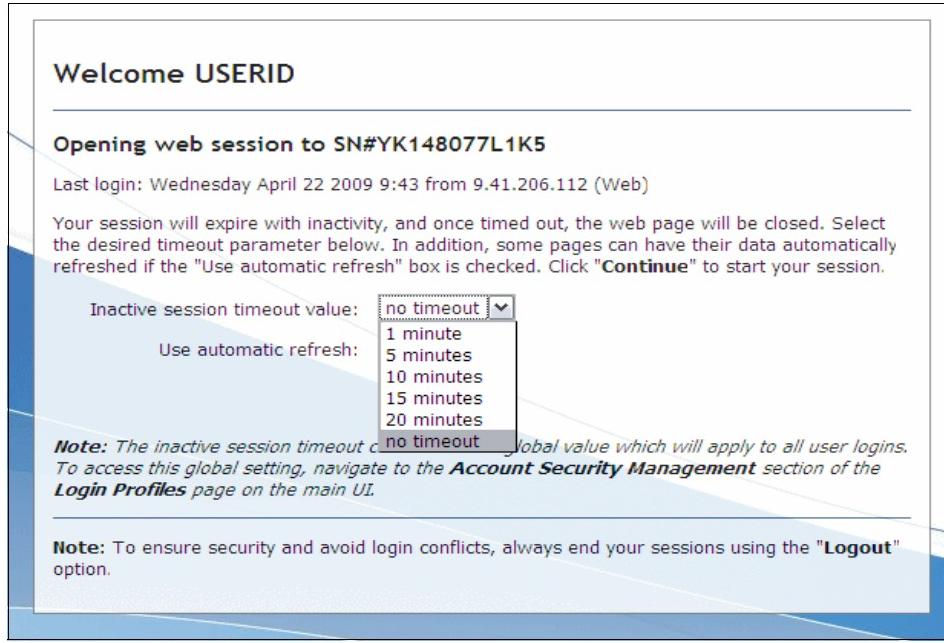


Figure A-3 Select time-out parameter

3. After successful login you will see the status page of the AMM. This page gives a short overview of the health of the chassis and the blades. Click **Remote Control** in the menu under Blade Tasks, as shown in Figure A-4 on page 499. Verify that there is no remote control session in progress by observing the remote control status. The Refresh button allows to refresh the status. Then scroll down to Start Remote Control.

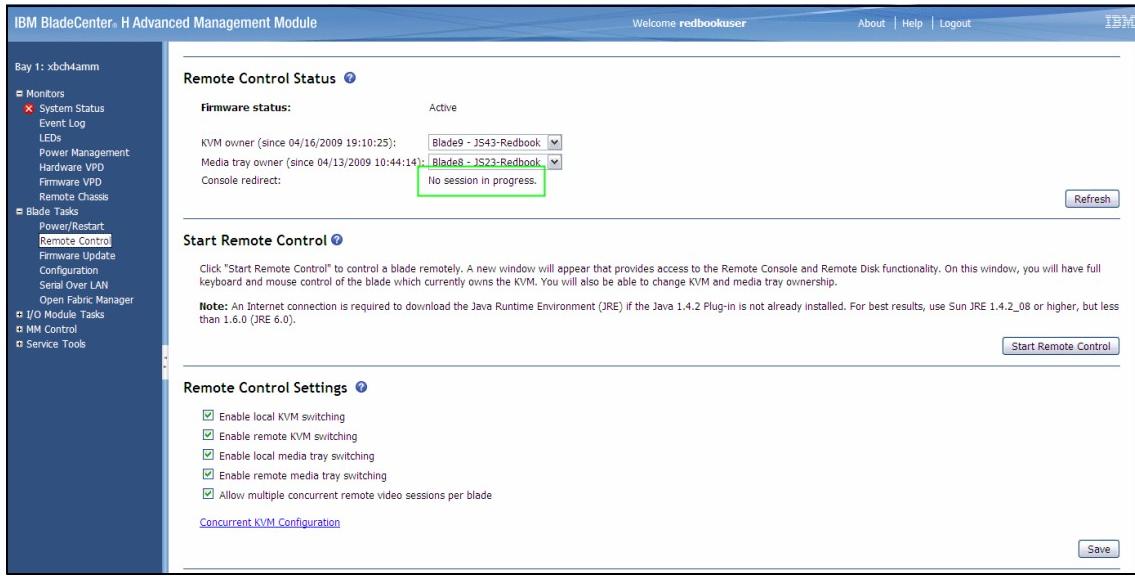


Figure A-4 Blade Remote Control options

- Click **Start Remote Control** as shown in Figure A-5. A new window will open with the remote control Java applet. Be sure that there are no popup blockers running or configure them to allow the popup windows from the AMM. It may take some time for the window to appear and load the applet.

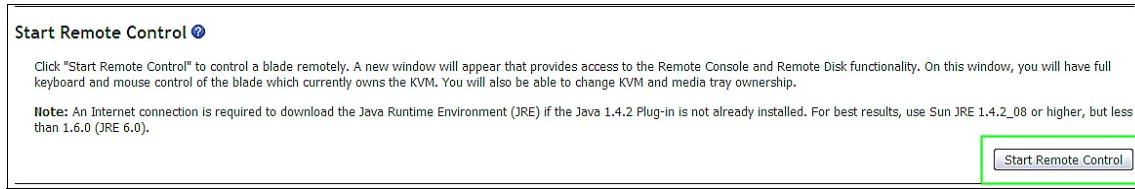


Figure A-5 Start remote control

The remote control Java applet will start in a new window. Figure A-6 shows remote control with remote media and remote console. Use the KVM drop-down list to switch the console between the available blades. The red A shows that this is an analog video session. As mentioned earlier, there can only be one analog session active. The physical console shows the same view as the remote console. Concurrent KVM will be shown only via remote control and is represented by a red D. JS23/JS43 does not support cKVM.



Figure A-6 Remote control - remote console and remote disk

Serial Over LAN

Serial over LAN (SOL) provides a means to manage servers remotely by using a command-line interface (CLI) over a Telnet or secure shell (SSH) connection. SOL is required to manage servers that do not have KVM support. SOL provides console redirection for both BIOS and the blade server operating system. The SOL feature redirects server serial-connection data over a LAN without the need for special cabling. The SOL connection enables blade servers to be managed from any remote location with network access. The advantages of SOL include:

- ▶ Remote administration without keyboard, video, or mouse (headless servers)
- ▶ Reduced cabling and no need for a serial concentrator
- ▶ Standard Telnet interface that eliminates the need for special client software

The IBM BladeCenter management module command-line interfaces provide access to the text-console command prompt on each blade server through an SOL connection, enabling the blade servers to be managed from a remote location.

In the BladeCenter environment, the integrated system management processor (ISMP) and network interface controller (NIC) on each blade server route the serial data from the blade server serial communications port to the network infrastructure of the BladeCenter unit, including an Ethernet-compatible I/O module that supports SOL communication. BladeCenter components are configured for SOL operation through the BladeCenter management module. The management module also acts as a proxy in the network infrastructure to couple a client running a Telnet or SSH session with the management module to an SOL session running on a blade server, enabling the Telnet or SSH client to interact with the serial port of the blade server over the network.

Because all SOL traffic is controlled by and routed through the management module, administrators can segregate the management traffic for the BladeCenter unit from the data traffic of the blade servers. To start an SOL connection with a blade server, first start a Telnet command-line interface session with the management module. When this Telnet or SSH command-line interface session is running, you can start a remote-console SOL session with any blade server in the BladeCenter unit that is set up and enabled for SOL operation.

You can establish up to 20 separate Web-interface, Telnet, or SSH sessions with a BladeCenter management module. For a BladeCenter unit, this enables you to have 14 simultaneous SOL sessions active (one for each of up to 14 blade servers) with 6 additional command-line interface sessions available for BladeCenter unit management.

With a BladeCenter S unit you have 6 simultaneous SOL sessions active (one for each of up to 6 blade servers) with 14 additional command-line interface sessions available for BladeCenter unit management. If security is a concern, you can use Secure Shell (SSH) sessions, or connections made through the serial management port that is available on the Advanced Management Module, to establish secure Telnet command-line interface sessions with the BladeCenter management module before starting an SOL console redirect session with a blade server.

SOL has the following requirements:

- ▶ An SOL-capable blade server such as the JS23/JS43.
- ▶ An Ethernet switch module or Intelligent Pass Through module is installed in Bay 1 of a BladeCenter.
- ▶ SOL is enabled for those blades that you wish to connect to via SOL.
- ▶ The Ethernet switch module must be set up correctly.

For details about how to set up SOL, see the Serial over LAN setup Guide, which can be found at:

<http://www-304.ibm.com/systems/support/supportsite.wss/docdisplay?lnocid=MIGR-54666&brandind=5000020>.

The System Management Service menu section contains an example of how to establish a Telnet or SSH connection to the management module and then an SOL console. See “Power on the IBM BladeCenter JS23 and JS43 using a Telnet or SSH session into the Advanced Management Module” on page 503.

System Management Services menu

The System Management Services Menu (SMS menu) is considered the main configuration interface of the IBM BladeCenter JS23/JS43. It is described in *Installation and User's Guide for IBM BladeCenter JS23/JS43*, (Type 7778-23X). This guide is delivered on CD with each JS23/JS43.

One of the first settings you may wish to change is the console that you want to use. When the blade starts the first time it comes up using the graphical console

as the active console if you do not select the SOL console as the active console. The SOL session cannot be used at this time to access the SMS menu to perform configuration tasks. To switch from the physical console to an SOL console you have to enter the SMS menu over the physical console or Remote Control. See “Graphical console” on page 494 about available consoles and how to use them.

To enter the SMS menu the blade has to go through the POST. You have to power on the blade or make a restart to be able to enter the SMS menu. As mentioned before, the SMS menu will only be available on the active console. The nonactive console can only be used to access the operating system.

Power on the IBM BladeCenter JS23 or JS43 using the Advanced Management Module Web interface

To power on or restart the blade you may use the Advanced Management Module (AMM) Web interface or a Telnet or SSH session to your AMM.

1. Log on to your AMM Web interface. The default IP address is 192.168.70.125. The default account is USERID with password PASSW0RD.

Note: Remember that the 0 in PASSW0RD is a zero.

2. In the left menu, click **Power/Restart** under Blade Tasks as shown in Figure A-7.



Figure A-7 Power/Restart in the left menu of the AMM Web interface

The power status of the blade is visible.

3. Click the check box in front of the blade. Then choose the power action below the table with the blades. Use **Power On Blade** to start the blade or **Restart Blade** to restart the blade. See Figure A-8 on page 503.

Blade Power / Restart

Blade selection and status

Click the checkboxes in the first column to select one or more blades; then, click one of the actions in the action list below the table and click "Perform Action" to perform the desired action.

	Bay	Name	Pwr	Local Pwr Control	Wake on LAN	Console Redirect
<input type="checkbox"/>	1	SN#YL30W7310006	Off	Enabled	N/A	
<input checked="" type="checkbox"/>	2	JS23-Redbook	Off	Enabled	N/A	
<input type="checkbox"/>	3	1GB-2GB Mixed	On	Enabled	N/A	
<input type="checkbox"/>	4	SN#ZK12HJ65L13F	Off	Enabled	On	
<input type="checkbox"/>	5	Free	Off	Enabled	N/A	
6 No blade present						

Available actions

Power On Blade

Standard actions

- Power On Blade
- Power Off Blade
- Shut Down OS and Power Off Blade
- Restart Blade
- Restart Blade with NMI
- Enable Local Power Control
- Disable Local Power Control
- Enable Wake on LAN
- Disable Wake on LAN
- Restart Blade System Mgmt Processor

POWER specific actions

- Restart Blade and clear NVRAM
- Restart Blade with Diagnostic Boot
- Restart Blade with Diagnostic Boot and Default Bootlist

Perform action

Figure A-8 Power/Restart blade options

Note: The Restart Blade option will perform a power off and a power on of your selected blade. The operating system will not shut down properly. Use this option only when there is no operating system running or the blade is in POST, SMS, or Open Firmware prompt.

The blade will perform the requested action.

4. Refresh this Web page to see a status change. Now use the console of your choice to work with the blade. Consoles are described in “Consoles of the IBM BladeCenter JS23 and JS43” on page 494.

Power on the IBM BladeCenter JS23 and JS43 using a Telnet or SSH session into the Advanced Management Module

To start or power cycle the blade via a Telnet or SSH session, log on with a Telnet or SSH client on the AMM. The default IP address is 192.168.70.125. The

default account is USERID with password PASSW0RD. See Example A-1 on page 505.

Note: Remember that the 0 in PASSW0RD is a zero.

Help is available via the command **help** or **help {command}**. Every command may be executed with one of these options to show the online help for the command:

```
env -h  
env -help  
env ?
```

This example uses the command **env** to show available options to get help. The *Management Module Command-line Interface Reference Guide* you can find online at:

<http://www-304.ibm.com/systems/support/supportsite.wss/docdisplay?lnodeid=MIGR-54667&brandind=5000020>.

After logon it is advisable to change the time-out of the Telnet or SSH session with the command **telnetcfg**. The command will be issued to the current primary (active) management module.

To get a list of available targets in a BladeCenter chassis, use the command **list -l {number of levels}**. The output in Example A-1 on page 505 shows, at the beginning of the **list -l 2** command, that the first management module is the active one. The **telnetcfg** command uses this active AMM as target to extend the time-out of a Telnet session or switch it off. Use 0 to switch the time-out off or use any value between 1 and 4,294,967,295 seconds. During installation it is a good idea to switch the time-out off so that the console will not disconnect.

Every command that is executed has a target. This target is specified by the **-T** option. To make the commands that you work with shorter, especially when you have to work for a long time only with a single blade, the environment command **env** is helpful. Specify the target that will be used. In Example A-1 we used the **blade3** as target. The resulting command is **env -T blade[3]**.

The power state can be managed with the **power** command. To query the power status, use the **-state** parameter. Power the blade on or off with **-on** or **-off** and power cycle the blade with the **-cycle** parameter. The parameter **-c** establishes an SOL session as soon as the blade is SOL ready. As an alternative to the **-c** parameter, you may issue the **console** command to establish an SOL session.

Example: A-1 Use of the power command

```
login as: USERID
Using keyboard-interactive authentication.
password:

Hostname:          moon.ibm.com
Static IP address: 172.16.0.225
Burned-in MAC address: 00:14:5E:DF:AB:28
DHCP:             Disabled - Use static IP configuration.
Last login: Friday June 20 2008 17:37 from 9.3.4.254 (SSH)

system> list -l 2
system
    mm[1]      primary
    power[1]
    power[2]
    power[3]
    power[4]
    blower[1]
    blower[2]
    blower[3]
    blower[4]
    switch[1]
    switch[3]
    switch[4]
    blade[1]   SN#YL30W7310006
    blade[2]   JS23-Redbook
    blade[3]   1GB-2GB Mixed
    blade[4]   SN#ZK12HJ65L13F
mt[1]
    storage[1]
    storage[2]
system>
system> telnetcfg -t 0 -T mm[1]
OK
system> env -T blade[2]
OK
system:blade[2]> power -state
Off
system:blade[2]> power -on -c
Establishing an sol connection. This may take a few minutes.
```

You may exit from the SOL session and return to the Advanced Management Module CLI by using the key combination ESC+[. This key combination can be defined in the AMM Web Interface.

Choosing the active console after first power on

When the blade is powered on for the first time as described in “Power on the IBM BladeCenter JS23 or JS43 using the Advanced Management Module Web interface” on page 502 or in “Power on the IBM BladeCenter JS23 and JS43 using a Telnet or SSH session into the Advanced Management Module” on page 503 you will see—depending on the console you have chosen—a message to select this console as the active console.

Note: Ensure you select the active console session using the 0 in a relatively quick time frame. This session will time out and force you to reboot the blade to re-establish the console session.

Figure A-9 shows an SOL console and Figure A-10 on page 507 shows a picture from the graphics console using Remote Control.

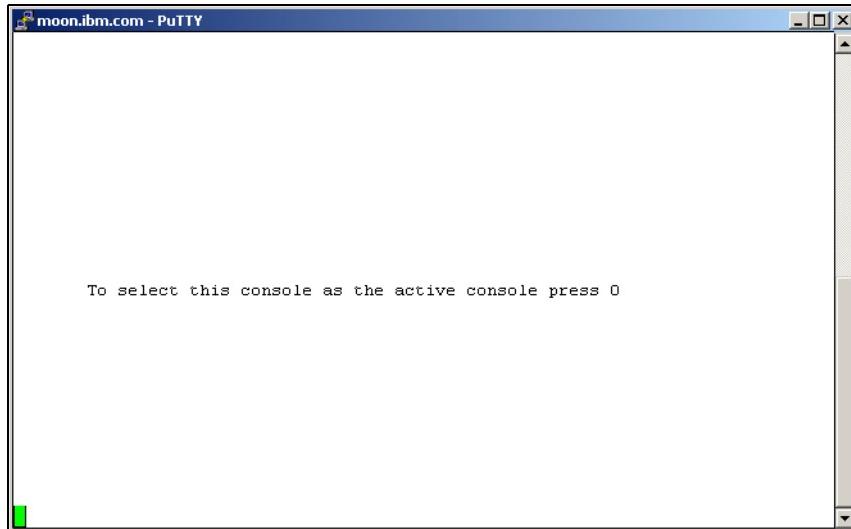


Figure A-9 Serial over a LAN console - select active console

Depending on the console you have open you have to enter a 0 to activate the SOL console or a 1 to activate the physical console. If you do not make your decision, the selection will default to SOL.

You need to enter the SMS menu over the physical console to change the active console, in this case as described in the next steps.



Figure A-10 Physical console shown with remote control - select active console

After a console is chosen as active console, either by the user or automatically, the system will show the Power On Self Test (POST).

IBM BladeCenter JS23 and JS43 Power On Self Test (POST)

As with previous JS2x blades, there are no System Reference Codes (SRC) shown on a console during POST. The System Reference Codes can be found in the Advanced Management Module under Blade Service Data in the Service Tools section. Click the blade in the table from which you wish to see the System Reference Codes. Click **System Reference Codes**. The table with the System Reference Codes can be refreshed by clicking **Refresh**.

The POST prints the words Memory, Keyboard, Network, SCSI, and Speaker as SMS loads. Press 1 to enter the SMS menu which is shown in Figure A-11 on page 508. This is the only response you see during power on or restart of the blade on an active console. During this printout you are able to choose one of the options from the menu. You may:

- ▶ (1) Enter the System Maintenance Services Menu
- ▶ (5) Use Default Boot List

- ▶ (6) Use Stored Boot list
- ▶ (8) Enter Open Firmware Prompt

The stored boot list used to load the operating system will be the default.

Click the number 1 to enter the SMS menu.

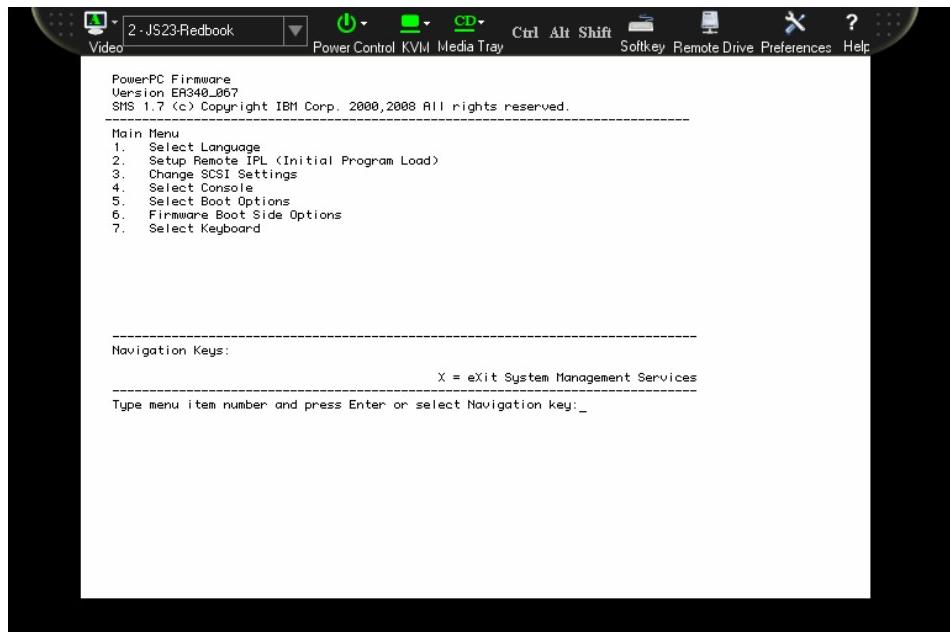


Figure A-11 JS23/JS43 SMS

Select the active console using the System Maintenance Services menu

When the blade is going through the POST, you can enter the System Maintenance Services menu.

To change the current active console in the SMS menu, click 5 to select the console. See Example A-2.

Example: A-2 SMS menu - main menu

```
Version EA330_031
SMS 1.7 (c) Copyright IBM Corp. 2000,2008 All rights reserved.

-----
Main Menu
1. Select Language
```

2. Setup Remote IPL (Initial Program Load)
3. Change SCSI Settings
4. Select Console
5. Select Boot Options
6. Firmware Boot Side Options

Navigation Keys:

X = eXit System Management Services

Type menu item number and press Enter or select Navigation key:

After you clicked **4** to select a console, you will see on the physical and on the SOL console a message notifying you that you have to enter 0 or 1 depending on the console you use. Figure A-10 on page 507 shows the message on an SOL console and Figure A-9 on page 506 shows the message on the physical console. Switch to the console you want to activate and enter the number as shown on this console. Either use 0 or 1. This console will be the next active console. If you choose to do nothing, the system will revert back to the original setting.

All other options are similar to other blades with Power CPU or System p servers.

Open Firmware interface

The Open Firmware prompt is a powerful interface into the system firmware of the JS23/JS43. Open Firmware is described in the IEEE 1275 standard. More general information about Open Firmware can be found at:

<http://www.firmworks.com>

The implementation of Open Firmware that is used by IBM on the JS23/JS43 to build the system firmware may have its own extensions and may also not implement each feature or function. The information on this page may not apply fully to the IBM BladeCenter JS23/JS43.

In certain situations it is helpful to use the Open Firmware prompt to change system settings or help troubleshoot problems. Also, the SMS menu is considered the main configuration interface. A few commands will be explained that help getting SAN booted with QLogic and Emulex host bus adapters configured. The commands can be used to perform the following tasks:

1. Identify FCode and firmware level.

2. Identify the World Wide Port Name and/or World Wide Node Name.
 3. Set the connection type.
 4. Set the transfer rate.
 5. Query available targets.

This appendix is split into a section about the QLogic host bus adapter and the Emulex host bus adapter. We start with a description of how to get access to the Open Firmware prompt.

Get access to the firmware prompt

Use a console of JS23/JS43 and power on or restart the blade. See “Graphical console” on page 494 or “Serial Over LAN” on page 500 about available consoles. “System Management Services menu” on page 501 describes how to power on a blade.

When the blade shows the POST, enter 8 to get access to the Open Firmware prompt. During POST the words Memory, Keyboard, Network, SCSI, and Speaker are written to the console, as shown in Example A-3. When this happens, enter 8 before the word speaker appears. The JS23/JS43 does not show the service request codes on the console, as JS20 or JS21 do. These SRCs are shown in the Web interface of the Advanced Management Module. The content of Example A-3 is the only information shown during POST.

Example: A-3 IBM BladeCenter JS23/JS43 POST

1 = SMS Menu

5 = Default Boot List

Memory	Keyboard	Network	SCSI	Speaker	ok
--------	----------	---------	------	---------	----

After entering the Open Firmware prompt, you see the command prompt shown in Example A-4.

Example: A-4 Open Firmware command prompt

0 >

Note: You may leave the System Maintenance Services Menu from the main menu with 0 to the Open Firmware prompt.

Boot settings are stored in the NVRAM of the system. The Open Firmware allows you to verify them with the **printenv** command. The command accepts an optional variable name to display only the content of the given variable. To display your boot device you may use the command **printenv boot-device** as shown in Example A-5. You see a iSCSI boot device entry.

ciaddr	iSCSI initiator IP address
giaddr	Gateway IP address
subnet-mask	Subnet mask of the initiator IP address
itname	IQN from the initiator
iport	iSCSI port
ilun	LUN to boot from on the iSCSI target
iname	IQN of the iSCSI target device
siaddr	iSCSI target portal IP address

Example: A-5 printenv boot-device output with iSCSI boot device

```
0 > printenv boot-device
----- Partition: common ----- Signature: 0x70 -----
boot-device
/pci@800000020000203/iscsi-toe@1,1:iscsi,ciaddr=172.16.1.44,giaddr=172.16.1.1,subnet-mask=255.255.255.0,itname=iqn.2008-06.com.ibm.blade5port1,iport=3260,ilun=0,iname=iqn.2008-06.com.ibm:storage1.disk1,siaddr=172.16.1.195,2
ok
0 >
```

The **show-devs** command is used to display all devices in a device path known by the system firmware. A part of the device path might be given as parameter to show only child devices from this device path. Without a given device path the command will return a listing of all known devices. To figure out the device path of the fibre channel host bus adapter the full device tree is required. Enter the

command show-devs on the Open Firmware prompt as shown in Example A-6. The output of the command is shortened to show only the important part of information for the explanation in this section.

Example: A-6 show-devs example output

```
0 > show-devs
00000208dda0: /ibm,serial
00000208eb98: /chosen
.
.
.
0000022fe48: /pci@800000020000202
0000022acb78: /display@1
00000223a0f8: /pci@800000020000204
0000022affe0: /fibre-channel@0
0000022c3da0: /disk
0000022c4790: /fibre-channel@0,1
0000022d8550: /disk
000002244440: /pci@800000020000205
0000022d8f40: /pci@0
0000022dae0: /ethernet@4
0000022eb198: /ethernet@4,1
ok
0 >
```

Look in the device path for the fiber-channel adapters. In a JS23/JS43 there are always dual port host bus adapters used, represented by the following two entries:

Port 1 /pci@800000020000204/fibre-channel@0
Port 2 /pci@800000020000204/fibre-channel@0,1

This information will be required in the next steps to change the topology and the link speed of the host bus adapter.

QLogic host bus adapter

This section describes how to

1. Retrieve the World Wide Node Name.
2. Identify FCode and the firmware level.
3. Set the transfer rate.
4. Set the connection mode.
5. Query available targets and LUNs.

The examples in this section were created using a CFFh combo card with the firmware 4.00.24 and FCode 1.25.

Identify your fiber channel host bus adapter as described in Example A-6 on page 512. The device tree in your system may differ from the example shown here. With this information you can build the command to select the device. Enter the command:

```
" /pci@800000020000204/fibre-channel@0" select-dev
```

to select the first host adapter port. The second HBA port is selected with the command:

```
" /pci@800000020000204/fibre-channel@0,1" select-dev
```

Note that there is a space between the first “ and /. Example A-7 shows the output of this command. The link of the adapter port will come up and the adapter will log in to the switch. You can now verify the name server of a connected SAN Switch about the World Wide Node and Port Name.

Example: A-7 Select the fiber channel host bus adapter port

```
0 > " /pci@800000020000204/fibre-channel@0" select-dev QLogic QM13472 Host  
Adapter Driver(IBM): 1.25 12/11/06  
Wait for link up - |/\|/\|/  
Firmware version 4.00.24  
ok  
0 >
```

To identify the World Wide Node Name you may use the Advance Management module. Alternatively, you can use the command **my-wwn** to retrieve the World Wide Port Name of the host adapter port. To display the retrieved World Wide Port Name enter a dot followed by Enter. Example A-8 shows the output of this command.

Example: A-8 Display World Wide Port Name of a QLogic host bus adapter port

```
0 > my-wwn ok  
1 > . 2100001b32005216 ok  
0 >
```

Firmware version and FCode level of the HBA can be shown with the command **version**. Example A-9 shows the output of this command. The FCode version is on the current level at the time of this writing, but the firmware version can be upgraded to 4.00.27.

Example: A-9 Firmware and FCode versions

```
0 > version QLogic QM13472 Host Adapter Driver(IBM): 1.25 12/11/06
```

```
Firmware version 4.00.24
ok
0 >
```

In case of the usage of an Optical Pass Through Module, it is necessary to change the transfer rate that is set, per default, to Auto Negotiation on the 4 GB host bus adapter to a fixed value of 2 GB. The Optical Pass Through Module can only handle transfer rates up to 2 GB. Auto Negotiation will not work with 4 GB host bus adapters. To change the transfer rate, verify the current settings of the HBA first. Use the command **show-settings** as shown in Example A-10 on page 514.

Example: A-10 Settings of the QLogic HBA

```
0 > show-settings QLogic QM13472 Host Adapter Driver(IBM): 1.25 12/11/06
Firmware version 4.00.24
Serial#
node-wwn 2001001b 32205216
port-wwn 2101001b 32205216
Current HBA Connection Mode: 2 - Loop preferred, otherwise point-to-point
Current HBA Data Rate: Auto-negotiated
Current Fcode Mode: qlc
ok
0 >
```

The **show-settings** command also shows the firmware and FCode versions as well as node and port names. Data rate is shown as well. The adapter is currently configured for Auto Negotiation. To change the data rate to a fixed data rate of 2 Gb, use the command **set-data-rate** as shown in Example A-11.

Example: A-11 Change the data rate of a QLogic HBA

```
0 > set-data-rate Current HBA Data Rate: Auto-negotiated
Do you want to change it? (y/n)
Choose HBA Data Rate:
0 - 1 Gigabit
1 - 2 Gigabit
2 - Auto-negotiated
3 - 4 Gigabit
enter: 1Current HBA Data Rate: 2 Gigabit
ok
0 >
```

Example A-11 shows that the data rate is changed from auto negotiation to a fixed value of 2 Gb. Enter the number in front of the value that you want defined here. The example uses 1. The changed value will be returned.

Depending on your fiber channel targets and the connectivity that you use to connect to them, you may wish to change the connection type to loop or to point-to-point. Use the command **set-connection-mode** to do the change, as shown in Example A-12. The command returns the current setting and lets you change to a new one. The possible options are shown. Select the corresponding number and click Enter.

Example: A-12 Change connection mode of an QLogic HBA

```
0 > set-connection-mode Current HBA Connection Mode: 2 - Loop preferred, otherwise
point-to-point
Do you want to change it? (y/n)
Choose HBA Connection Mode:
0 - Loop Only
1 - Point-to-point only
2 - Loop preferred, otherwise point-to-point
enter: 1Current HBA Connection Mode: 1 - Point-to-point only
ok
0 >
```

The last command that we describe for the QLogic host bus adapter is the **show-children** command, which can be used to show available targets and LUNs to the HBA. Example A-13 shows output from one of our Virtual I/O servers. A DS4800 with 7 LUNs is shown. The DS4800 is connected to port 0 of the Switch and JS23/JS43 to port 3.

Example: A-13 List all available devices on the selected QLogic HBA

```
0 > show-children Adapter portID - 610300
***** Fabric Attached Devices *****
Dev# 0 (0 )  PortID 10000    Port WWN 203200a0b811a662
      LUN 2   DISK  IBM     1815    FASST 0914
      LUN 3   DISK  IBM     1815    FASST 0914
      LUN 4   DISK  IBM     1815    FASST 0914
      LUN 5   DISK  IBM     1815    FASST 0914
      LUN 6   DISK  IBM     1815    FASST 0914
      LUN 7   DISK  IBM     1815    FASST 0914
      LUN 8   DISK  IBM     1815    FASST 0914

ok
0 >
```

Remember that the described commands require that you have an HBA port selected and that they have effect only on the selected HBA port. You need to perform the necessary actions on both HBA ports.

To leave the Open Firmware prompt and restart the blade, use the command **reset-all**.

When no changes are made, the boot process can be started by leaving the Open Firmware prompt with the commands as shown in Example A-14.

Example: A-14 Leave Open Firmware prompt

```
1 > dev /packages/gui  
1 > obe
```

Emulex host bus adapter

This section describes how to

1. Retrieve the World Wide Node Name.
2. Identify the FCode level.
3. Set the link speed.
4. Set the connection mode.

The examples in this section were created using an Emulex CFFv with the FCode 3.10.a0.

Identify your fiber channel host bus adapter as described in Example A-6 on page 512. The device tree in your system may differ from the example shown here. With this information you can build the command to select the device. Enter the command

```
“ /pci@800000020000203/fibre-channel@0” select-dev
```

to select the first host adapter port. The second HBA port is selected with the command:

```
“ /pci@800000020000203/fibre-channel@0,1” select-dev
```

Note that there is a space between the leading “ and /. Example A-15 shows the output of this command. The link of the adapter port will come up and the adapter will log in to the switch. You are able now to verify the name server of a connected SAN Switch about the World Wide Node and Port Name.

Example: A-15 Select fiber channel port

```
0 > “ /pci@800000020000203/fibre-channel@0” select-dev ok  
0 >
```

World Wide Node and Port Name of the HBA port are shown by the command **host-wwpn/wwnn**. The same information can be retrieved via the Advanced Management Module Web interface under Hardware VPD. Example A-16 shows the WWPN and WWNN of the first port of an Emulex CFFv HBA.

Example: A-16 Display the World Wide Node and Port Name of an Emulex CFFv HBA

```
0 > host-wwpn/wwnn Host_WWPN 10000000 c9660936
Host_WNN 20000000 c9660936
ok
0 >
```

The installed FCode level on the HBA can be shown with the command **check-vpd** or **.fcode**, as shown in Example A-17.

Example: A-17 Display FCode version of an Emulex CFFv HBA

```
0 > check-vpd
!!! LP1105-BCv Fcode, Copyright (c) 2000-2008 Emulex !!! Version 3.10a0

ok
0 >
0> .fcode Fcode driver version 3.10a0
ok
0>
```

To display the current link speed, use the command **link-speed** as shown in Example A-18. The command does not return the current link speed setting.

Example: A-18 Display actual link speed of an Emulex CFFv HBA

```
0 > link-speed
Current ....

Link Speed -- 2 Gb/s ok
0 >
```

Link speed or data rate of the Emulex HBA can be set with the command **set-link-speed**. The command will show the current link speed and the current setting, as shown in Example A-19.

Example: A-19 Set link speed of an Emulex CFFv HBA

```
0 > set-link-speed
Current ....

Link Speed -- 2 Gb/s
Link Speed Selected -- Auto Select
```

0. Auto Select Link Speed (Default)
1. 1 Gb/s Link Speed -- Only
2. 2 Gb/s Link Speed -- Only

4. 4 Gb/s Link Speed -- Only

Enter <x> to QUIT

Enter a Selection:

Enter the number of your choice and click Enter as shown in Example A-20. The NVRAM of the HBA will be updated.

Example: A-20 Changed link speed in NVRAM of the Emulex CFFv HBA

Enter a Selection: 2

```
Flash data structure updated.  
Signature          4e45504f  
Valid_flag        0000004a  
Host_did          00000000  
Enable_flag       00000005  
SFS_Support       00000000  
Topology_flag     00000000  
Link_Speed_flag   00000002  
Diag_Switch       00000000  
POST-Linkup       00000000  
Boot_id           00000000  
Lnk_timer         0000000f  
Plogi-timer       00000000  
LUN (1 byte)      00000000  
DID               00000000  
WWPN              0000.0000.0000.0000  
LUN (8 bytes)     0000.0000.0000.0000
```

```
*** Type reset-all to update. ***  
ok  
0 >
```

The connection type or topology setting can be shown with the command **.topology**. This command ready the NVRAM of the Emulex HBA and displays the value as human-readable text. To change the topology setting use one of the available commands, such as **set-auto-fcal**, **set-auto-ptp**, **set-fc-al** or **set-ptp**. **set-auto-fcal** sets loop preferred, otherwise point-to-point. **set-auto-ptp** sets point-to-point, otherwise loop. The two other commands set the connection type to point-to-point or loop. The default setting can be defined with **set-default-mode**. A restart of the blade is required. In Example A-21 you

can see that the topology is set to Point to Point. The set commands return nothing.

Example: A-21 Display connection topology of an Emulex CFFv HBA

```
1 > .topology Point to Point - Current Mode  
Manual Topology  
ok  
1 >
```

Remember that the described commands require that you have an HBA port selected and that they only have effect on the selected HBA port. You need to perform the necessary actions on both HBA ports.

To leave the Open Firmware prompt and restart the blade, use the command **reset-all**.

When no changes are made, the boot process can be started by leaving the Open Firmware prompt with the commands shown in Example A-22.

Example: A-22 Leave Open Firmware prompt

```
1 > dev /packages/gui  
1 > obe
```



B

SUSE Linux Enterprise Server AutoYaST

This appendix describes the SUSE AutoYaST tool to perform automated installations of SUSE Linux Enterprise Server 11.

We discuss the following topics:

- ▶ “AutoYaST introduction” on page 522
- ▶ “AutoYaST profile creation methods” on page 522
- ▶ “Create an AutoYaST profile using YaST Control Center” on page 522

AutoYaST introduction

The AutoYaST configuration tool allows a system administrator to install SUSE Linux Enterprise Server (SLES) on a large number of systems in parallel using an automated process. The AutoYaST profile is a file written using the Extensible Markup Language (XML). It contains responses to all the system configuration questions typically asked during a manual installation. This file is configurable to accommodate the installation of systems with homogeneous and heterogeneous hardware.

This appendix covers AutoYaST profile creation for SUSE Linux Enterprise Server 11.

Note: The procedures covered here are not applicable for previous SLES releases because the XML layouts are different.

AutoYaST profile creation methods

There are three methods to create an AutoYaST profile for systems using *identical* hardware:

1. Clone the install configuration information from a reference machine that was installed manually. This is covered in 9.4, “Linux network installation (detailed)” on page 374.
2. Use the YaST Control Center to create and modify the AutoYaST profile.
3. Use an XML editor to create a AutoYaST profile from scratch.

We cover the steps for option 2 here.

Create an AutoYaST profile using YaST Control Center

In this section, we use the AutoYaST configuration tool available in the YaST Control Center with a graphical interface to demonstrate how to create a basic XML file.

Note: This YaST tool can run in graphical or text mode. A mouse can navigate through the graphical version of the tool while the text mode version requires Tab, Enter, Up/Down Arrow, and Space bar keys to navigate. Otherwise, there is no difference between the two modes and the same configuration options in both will result in the same XML file.

There are a lot of optional settings, but some are mandatory settings or dependencies. It is impossible to cover every possible configuration option, so we try to provide a general overview to make you familiar enough with the tool to navigate on your own.

Starting the YaST graphical interface

1. SSH as root with X11 forwarding enabled into a system running SLES 11. For example:

```
ssh -X root@9.3.20.18
```

Restriction: It is necessary to issue this command on a system with a running X Server to use the graphical interface.

2. From the command line type:

```
yast2
```

and a new window will appear with YaST Control Center at the very top as shown in Figure B-1 on page 524.

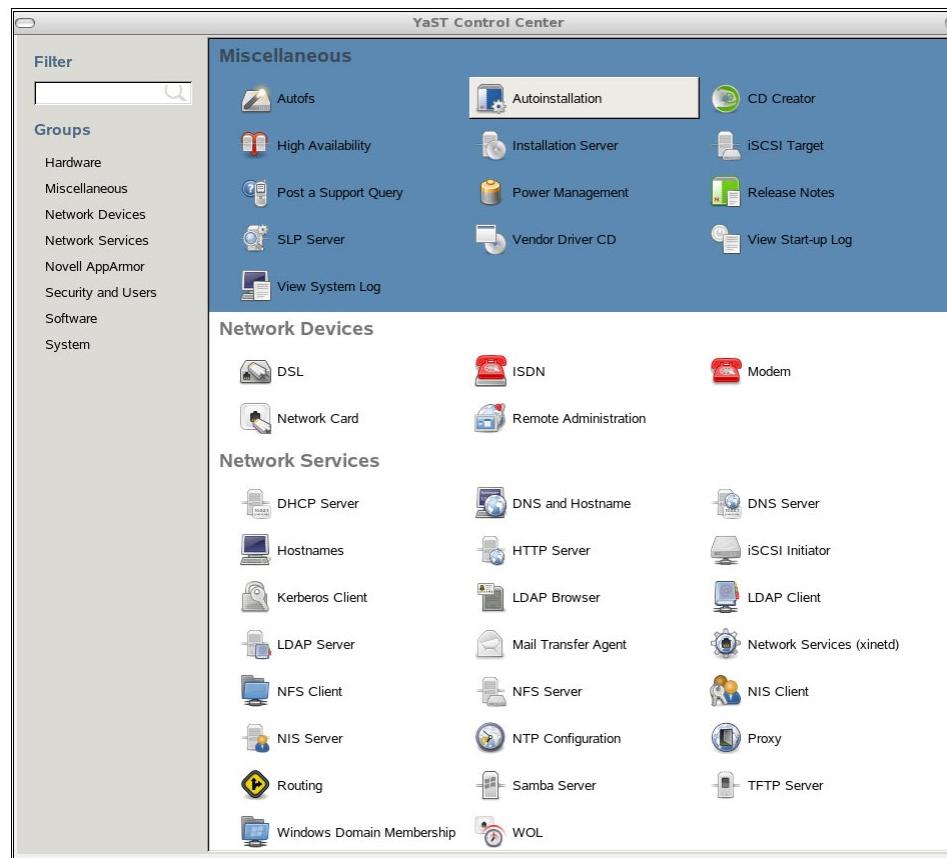


Figure B-1 YaST Control Center in graphics mode



Figure B-2 YaST Control Center in text mode

Navigating the YaST graphical interface

1. Start the YaST application, which opens a window as shown in Figure B-3.
Launch the **Autoinstallation** applet from the Miscellaneous section of YaST.

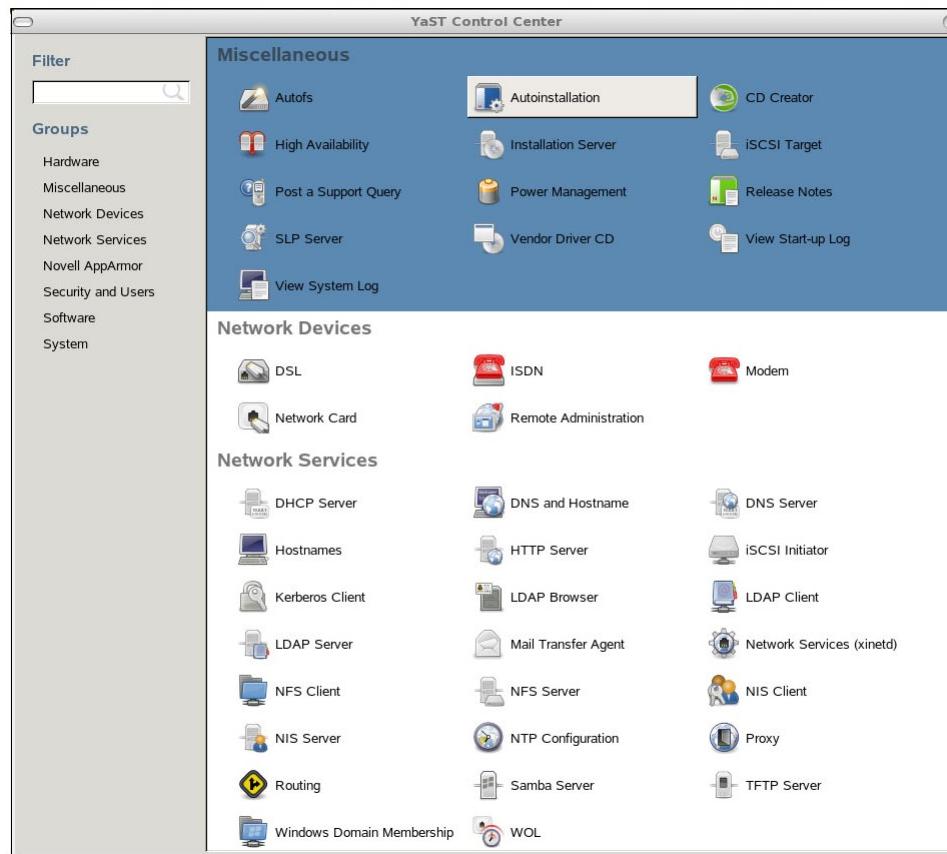


Figure B-3 Selecting the Autoinstallation option

2. After the selection, the main AutoYaST configuration window opens as shown in Figure B-4.

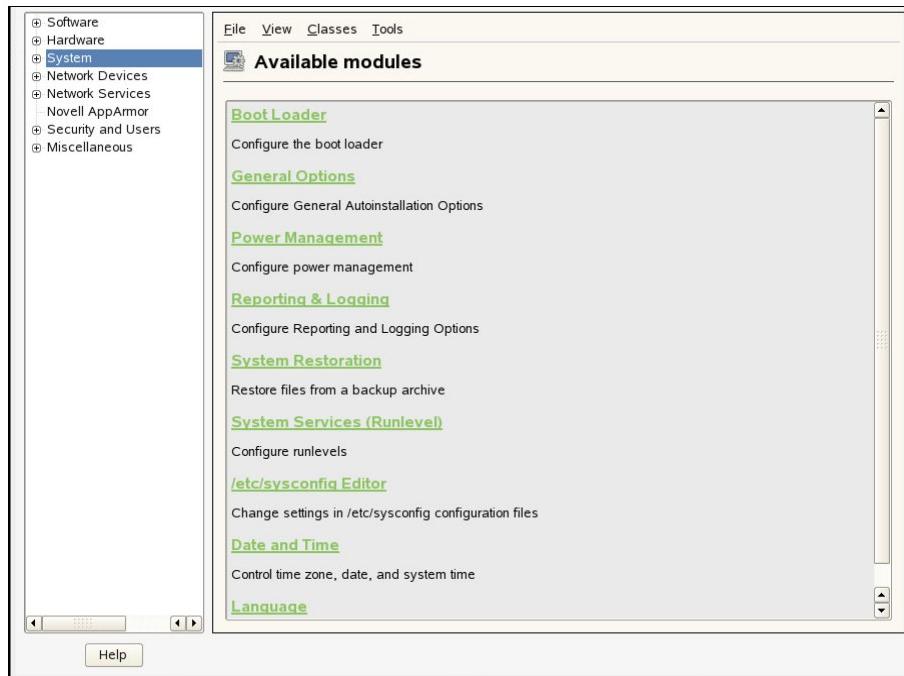


Figure B-4 Main AutoYaST menu (SLES 11)

3. Clone the configuration of the installation server by selecting **Tools** → **Create Reference Profile**, as shown in Figure B-5.

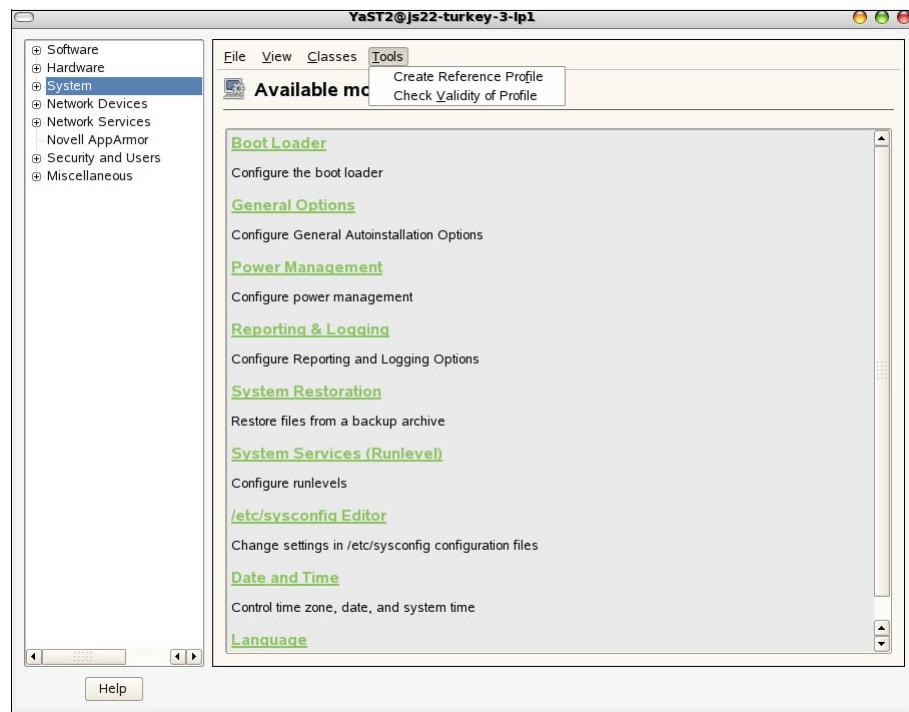


Figure B-5 Create a reference profile

4. A second window opens, as shown in Figure B-6. In addition to the default resources such as boot loader, partitioning, and software selection, it is possible to add other aspects of your system to the profile by checking items in the Select Additional Resources section. When ready, click **Create** so YaST can collect the system information and create the AutoYaST profile.

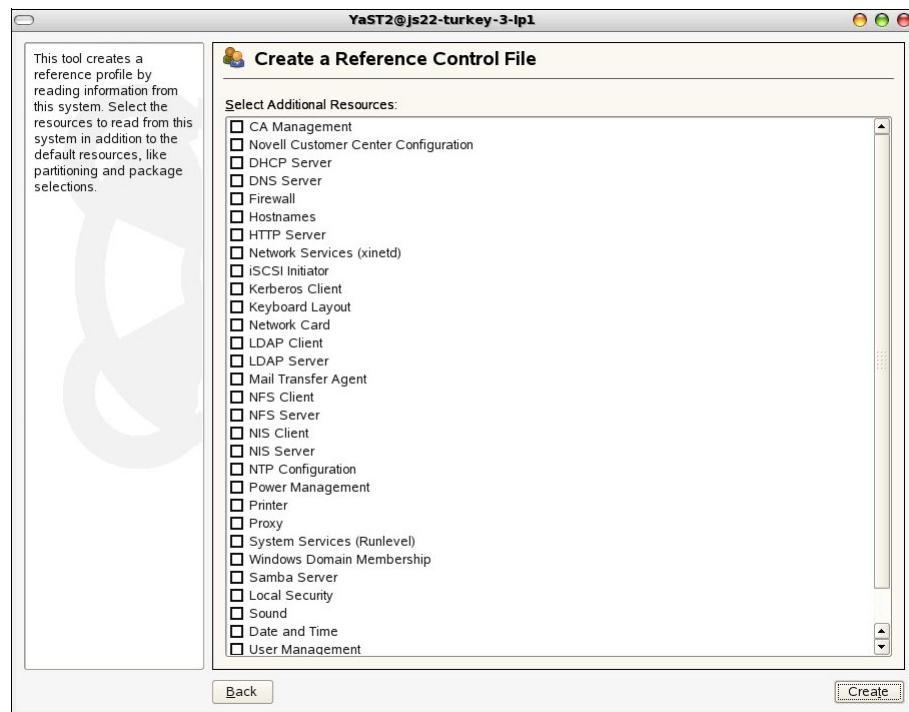


Figure B-6 Selecting additional resources

5. If the profile is complete and meets your requirements, select **File → Save** and enter a file name such as

```
sles11_autoinst.xml
```

Or adjust some of the options provided on the left side of the AutoYaST main menu option. Each option is given here with some highlights:

- a. Software - Options to select and configure the Online Update and Package Selection sections. Figure B-7 on page 530 shows the **Software → Package Selection** window. We chose the minimum software configuration to save time during installation.

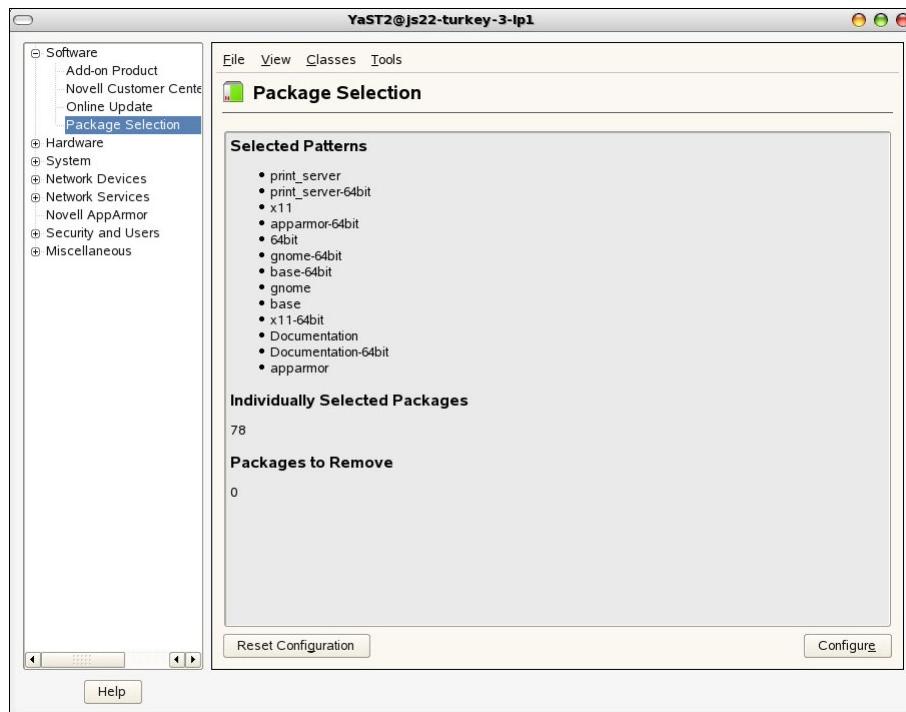


Figure B-7 AutoYaST software selection

- b. Hardware - Configures Partitioning, Sound, Printer, and Graphics Card and Monitor, if necessary. The Partitioning settings are critical for this configuration to work, so verify that they match your hard disk environment and that each partition meets the minimum SuSE partition size requirements.
- c. System - Sets the general system information such as language configuration, time zone, other locale-related settings, logging, and run-level information in this option. The most important configuration is the Boot Loader configuration, which is cloned from the /etc/lilo.conf file of the system running the YaST tool so it may require some editing. Edit it by pressing **Configure** → **Edit**. Click **Finish** when you are done.
- d. Network Devices - Sets the network adapter information. You can set network module information and IP details here.
 - i. Click **Configure** to open the menu.
 - ii. The Network cards configuration main menu opens. Select **Configure** to add interfaces.
 - iii. Select **Traditional Method with ifup** → **Next**.

- iv. Remove any static IP configurations on the next panel and press **Add**. Some selections are already configured, such as Device Type: Ethernet. Type, for example, ehea, as module name for the adapter and click **Next**.
- v. In the Host name and name server section, choose **DHCP** for the Hostname and Domain Name (Global) and also choose **DHCP** for Name servers and the domain search list.
- vi. Click **OK** → **Next**. Interface eth0 is ready now.

To create interface eth1, repeat the steps. However, the DHCP settings are automatically taken from the eth0 configuration and the interface name automatically changes to eth1.

- e. Network Services - Configures network clients and daemons using this option. There are more than 15 daemons to choose from and all are optional.
- f. Security and Users - Creates users and configures security policies. To allow SSH through the enabled firewall, follow these steps:
 - i. Click **Firewall**.
 - ii. Select **Start-Up** and change the firewall setting to **When Booting**.
 - iii. Select **Interfaces** and change eth0 to **External Zone**.
 - iv. Change eth1 to **Internal Zone**.
 - v. In the Allowed Services section, select **Secure Shell** to enable Port 22 connections through the firewall. Disable **Protect from Internal Zone**.

To set the password requirements:

- i. Select **Local Security** from the Security and Users menu.
- ii. Click **Configure**.
- iii. Adapt the Security settings. The default is seven characters for the maximum password length.

It is mandatory to define the root user password to log in a root via SSH and the steps to set the password are:

- i. Select **User Management** from the left menu.
- ii. Click **Configure**.
- iii. Select **Users**, high-light **root** and its row, and change the Filter to **System Users** by clicking the **Set Filter** drop-down button as shown in Figure B-8 on page 532.

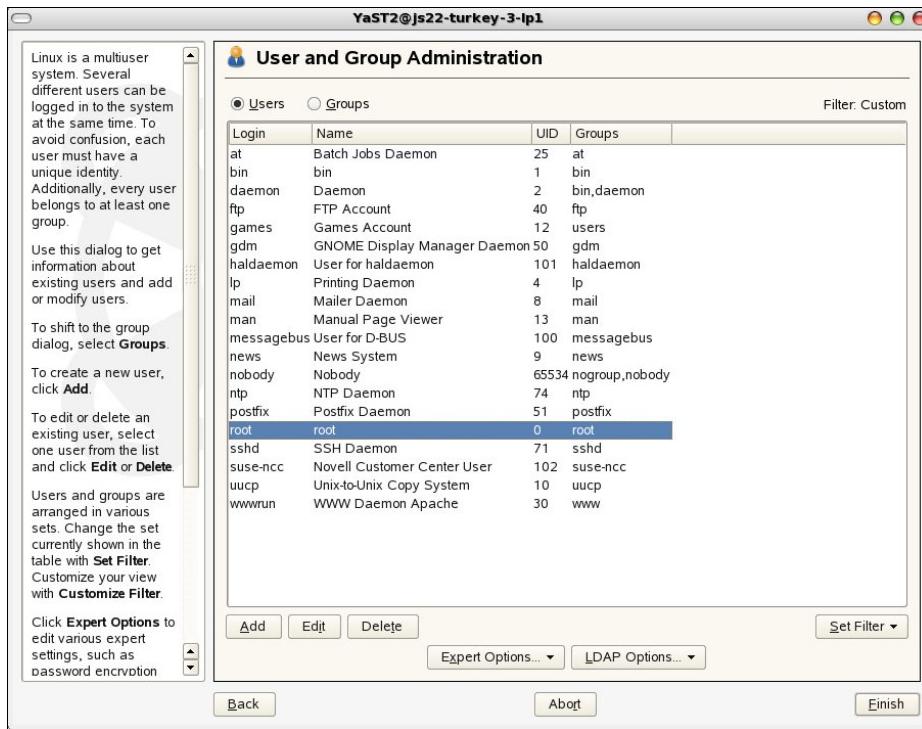


Figure B-8 Configure the root user

- iv. High-light **root** and its row again and press **Edit**.
 - v. Add the root user password. This password is saved encrypted in the XML file. Press **Accept** when finished.
 - vi. Click **Finish** to return to the AutoYaST main menu.
 - g. Misc - Allows you to add complete configuration files, or to add special scripts to run before and after the installation.
6. Remember to save the edits with **File → Save**.

Example: B-1 Part of newly created XML file

```
<?xml version="1.0"?>
<!DOCTYPE profile>
<profile xmlns="http://www.suse.com/1.0/yast2ns"
  xmlns:config="http://www.suse.com/1.0/configns">
  <add-on>
    <add_on_products config:type="list"/>
  </add-on>
  <bootloader>
```

```
<global>
  <activate>true</activate>
  <boot_chrp_custom>/dev/sda1</boot_chrp_custom>
  <default>Linux</default>
  <lines_cache_id>2</lines_cache_id>
  <timeout config:type="integer">80</timeout>
</global>
<uid>104</uid>
.
.
.
.
.
<user_password>*</user_password>
<username>pulse</username>
</user>
</users>
<x11>
  <color_depth config:type="integer">4</color_depth>
  <display_manager>gdm</display_manager>
  <enable_3d config:type="boolean">false</enable_3d>
  <monitor>
    <display>
      <max_hsync config:type="integer">42</max_hsync>
      <max_vsync config:type="integer">72</max_vsync>
      <min_hsync config:type="integer">30</min_hsync>
      <min_vsync config:type="integer">50</min_vsync>
    </display>
    <monitor_device>Unknown</monitor_device>
    <monitor_vendor>Unknown</monitor_vendor>
  </monitor>
  <resolution>640x480 (VGA)</resolution>
  <>window_manager>gnome</window_manager>
</x11>
</profile>
```



C

Additional Linux installation configuration options

This appendix describes some of the other options to install Linux natively or on an LPAR.

We cover the following configurations:

- ▶ “Basic preparations for a Linux network installation” on page 536
- ▶ “Virtual optical device setup and installation” on page 544

Basic preparations for a Linux network installation

This section provides all the basic information to set up services for a Linux network installation. In principle, this is not bound to a specific operating system or distribution that runs on the infrastructure server to provide the necessary services. Nevertheless, all descriptions in this section are based on general Linux services, commands, and parameters. We presume that the required files for all the services are already installed and that all the commands are issued with superuser rights.

Installing Linux using the network - General remarks

You always require the following services to perform a network installation:

- ▶ A running Bootstrap Protocol (BOOTP) service or a Dynamic Host Configuration Protocol (DHCP) service that includes BOOTP support to configure the network interface of a BladeCenter JS23.
- ▶ A running Trivial File Transfer Protocol (TFTP) service to serve the boot image to a BladeCenter JS23.
- ▶ It is necessary to set up *one* of the following services to provide the installation packages for a network installation after the boot image is loaded:
 - File Transfer Protocol (FTP)
 - Hypertext Transfer Protocol (HTTP)
 - Network File System (NFS)

This section shows how to configure DHCP, TFTP, and NFS services.

Note: If a firewall is running on the installation server, update the settings to allow traffic for the installation protocol.

Configuring a BOOTP or DHCP service

DHCP is an extension to the original BOOTP specification. As a result, you can use DHCP to provide the BOOTP information for booting using the network. The standard DHCP daemon is called *dhcpd*, but there are other DHCP daemons.

Note: The directory you use for the configuration files depends on the distribution. The following directories are possible examples:

- ▶ /etc/
- ▶ /etc/sysconfig/
- ▶ /etc/default/
- ▶ /etc/xinet.d/ (eXtended InterNET daemon configuration files)

The examples in this appendix use the most common directories. In general, the name of a configuration or script file is related to the name of the installed package. For example, if a DHCP daemon is called *dhcpd3-server*, you can find the configuration in */etc/dhcpd3-server.conf* and */etc/sysconfig/dhcpd3-server*, and the start/stop script is in */etc/init.d/dhcp3-server*.

The standard DHCP daemon is configured through two files: The */etc/sysconfig/dhcpd* file, which stores the basic configuration, and the */etc/dhcpd.conf* file, which contains the configuration information for each registered client. For a running service, the configuration actually used in most cases is copied in a subdirectory of */var/*. See Example C-1 for a simple client configuration stored in *dhcpd.conf*.

Note: Keep in mind that Example C-1 contains environment-specific Internet Protocol (IP) and Media Access Control (MAC) address information. One way to learn the MAC address of a JS23 BladeCenter is to use the BladeCenter management module. Select **Monitors** → **Hardware VPD** from the left-side options and then click your blade bay and then the **Ports** tab on the right side. For more information about how to customize *dhcpd.conf*, see the man pages of the *dhcpd* service and use **man dhcpd** on the command prompt.

Example: C-1 dhcpd.conf example

```
ddns-update-style none;
allow booting;
allow bootp;
always-reply-rfc1048 true;

shared-network TEST {
```

```
option routers 172.16.1.1;
subnet 172.16.1.0 netmask 255.255.255.0 {
    option broadcast-address 172.16.1.255;
    range dynamic-bootp 172.16.1.68 172.16.1.80;
    default-lease-time 444;
    next-server 172.16.1.197;
}

host JS23 {
    hardware ethernet 00:1a:64:44:21:53;
    fixed-address 172.16.1.79;
    filename "install";
}
}
```

You can find the start and stop scripts of Linux services in the /etc/init.d/ directory. To start the standard DHCP daemon, use the **/etc/init.d/dhcpd start** command. To restart the DHCP daemon, use the **/etc/init.d/dhcpd restart** command.

The filename “install”; lines in the dhcpcd.conf file point to a directory in the /tftpboot directory. We show how to configure it in the next section.

Tip for Linux beginners: The following tasks help you to double-check or troubleshoot a configuration in general.

- ▶ To trace messages of running services, type **tail -f -n 10 /var/log/messages** to get the last 10 messages and auto update if there are new messages.
- ▶ Connect to a running service with a local client, remote client, or both these clients and try to receive the data that you want.
- ▶ Make sure a changed configuration is activated by restarting a service *directly* after editing, for example:
 - a. **vi /etc/dhcpcd.conf**
 - b. **/etc/init.d/dhcpcd restart**

Configuring a Trivial File Transfer Protocol service

You can use the TFTP to provide a bootable image during a network installation. There are several implementations of TFTP daemons available. The standard TFTP daemon is called *tftpd*. In general, the *xinetd* or *inetd* super daemons are used to create a TFTP daemon. You can also run a TFTP daemon without one of the super daemons.

Install the *tftp-server rpm* if it was not installed by default, create a */tftpboot* directory if does not exist, and set the *tftp* service to *disable=no*. For example, on SLES11:

1. Type **yast** on the command line to start the YaST tool.
2. Scroll down to **Network Services**.
3. Select **TFTP Server** on the right side and press Enter.
4. In the new window, select **Enable** and create */tftpboot* in the Boot Image Directory section as shown in Figure C-1.

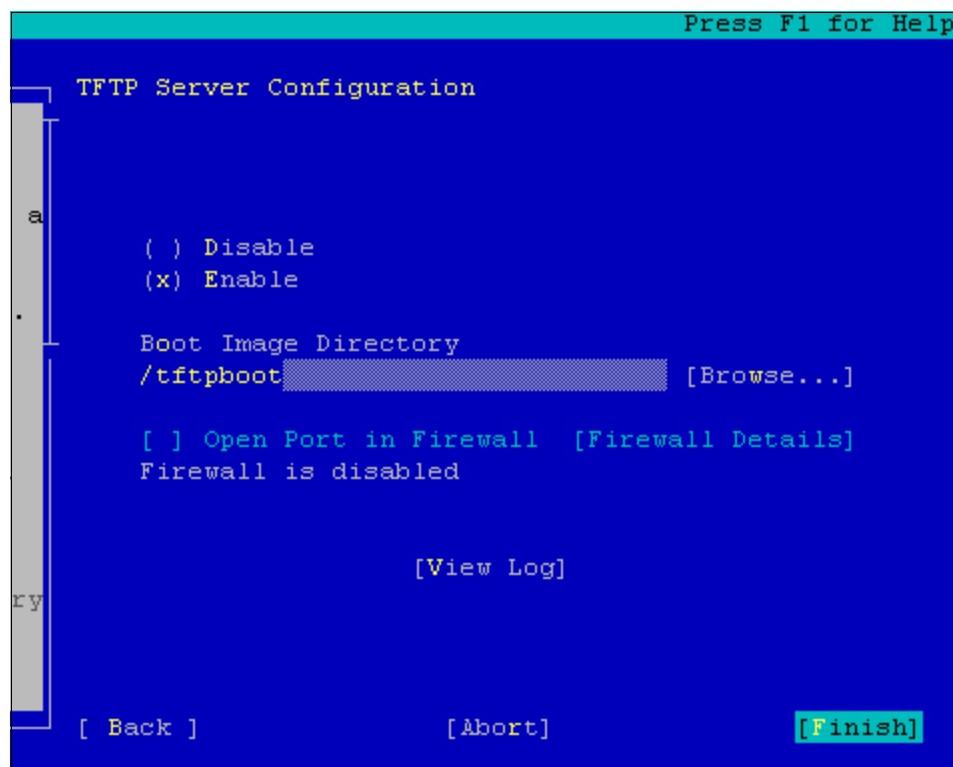


Figure C-1 Configure a TFTP server in SLES11

- Finally, scroll down to **[Finish]** and press the Enter key.

Example C-2 shows a TFTP daemon configuration for xinetd stored in /etc/xinet.d/tftpd.

Example: C-2 Configuring a TFTP daemon in the /etc/xinet.d/tftp file on SLES11

```
# default: off
# description: tftp service is provided primarily for booting or when a \
#               router need an upgrade. Most sites run this only on machines
#               acting as
#               "boot servers".
service tftp
{
    socket_type      = dgram
    protocol         = udp
    wait             = yes
    user             = root
    server           = /usr/sbin/in.tftpd
    server_args      = -s /tftpboot
    disable          = no
}
```

Tip: You can edit the etc/xinet.d/tftp file using a text editor as well.

Red Hat Enterprise Linux and SLES create TFTP from the xinetd daemon. Restart the xinetd service after the tftp configuration is complete by:

```
venus:/ # service xinetd restart
Shutting down xinetd: done
Starting INET services. (xinetd) done
```

The next step is to copy the bootable installation kernel into the /tftpboot directory. The bootable kernel must match the Linux distribution that you install. It must also match the *file name* listed in the dhcpcd.conf, as shown in Example C-1 on page 537. The process is slightly different for Red Hat Enterprise Linux and SLES, so we document both separately here.

Copying the SLES11 install kernel

To copy the SLES11 install kernel, use the following procedure:

- Mount the SLES11 DVD1 on the system running the tftp server. For example, on a system running SLES, type:

```
mount /dev/sr0 /mnt
```

2. Then enter

```
cp /mnt/suseboot/inst64 /tftpboot/install
```

Copying the Red Hat Enterprise Linux 5 install kernel

To copy the Red Hat Enterprise Linux 5 install kernel, use the following procedure:

1. Mount the Red Hat Enterprise Linux 5.2 DVD1 on the system running the tftp server. For example, on a system running Red Hat Enterprise Linux 5, type:

```
mount /dev/cdrom /mnt
```

2. Then enter

```
cp /mnt/images/netboot/ppc64.img /tftpboot/install
```

The next step is the preparation of the installation source directory and the corresponding service.

Configuring a Network File System Protocol service

With DHCP and TFTP protocols configured, you need to configure the installation source.

On SLES

This section shows how to set up a Network File System (NFS) server using the Installation Server utility provided in the YaST tool. Keep in mind that this is just one way to set up an NFS server. For this setup our installation server is running SLES 11. The configuration steps are:

1. Type **yast** on the command line.
2. In the YaST window, scroll down to **Miscellaneous** and select **Installation Server** on the right side.
3. Select **Configure as NFS Source** and enter the desired location source location directory. This example uses `/install` as shown in Figure C-2 on page 542.

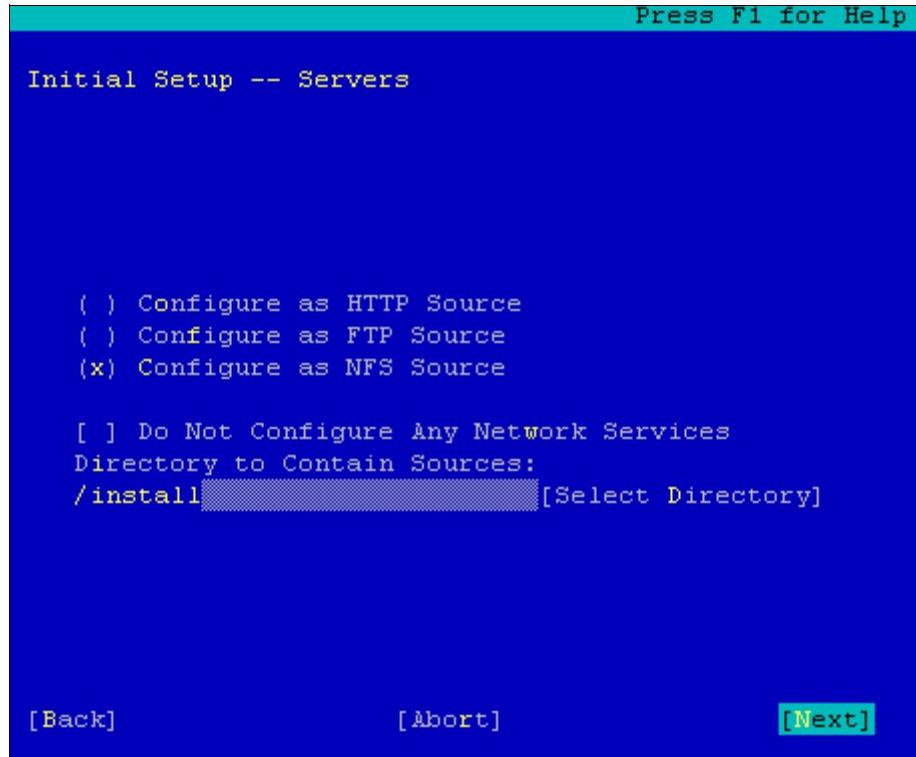


Figure C-2 Initial setup of SLES NFS installation server

4. Then click **[Next]**.
5. Leave the defaults for **Host Wild Card** and **Options**.
6. Click **[Next]**. With this, an NFS server serving /install is set up automatically.
7. Click **Add** to configure an installation source.
8. As Source Name, enter the desired name for this installation source, for example, sles11. This creates a subdirectory sles11 under /install.
9. At the Source Configuration window, if you have a DVD of the operating system, click the check box. Otherwise, if you have ISO images of the DVD contents, select the second check box and browse via Select Directory to the directory that contains all ISO images of all CDs (see Figure C-3 on page 543).
10. Click **[Next]** when finished.

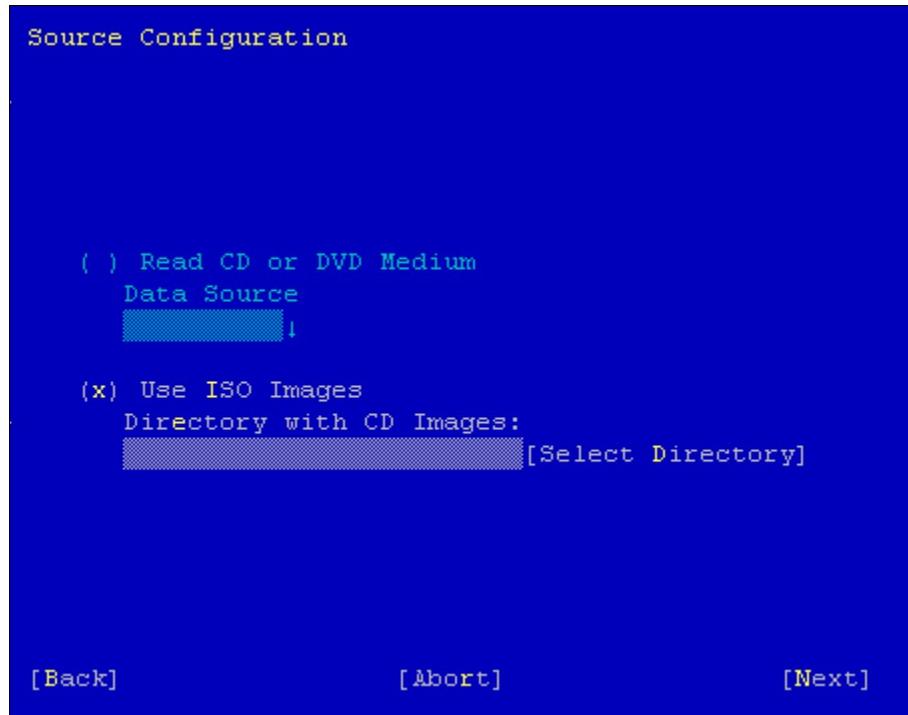


Figure C-3 Source configuration window

- 11.If you chose the Read CD or DVD Medium option given in Figure C-3, you will be prompted to insert the first DVD.
- 12.Insert SLES11 DVD1 into the BladeCenter media tray and press [**Continue**]. The data from DVD1 is copied to the /install/sles11/CD1 directory.

Note: If you used the CD option instead of a DVD, you will be prompted for the other CDs at this step.

- 13.Select [**Finish**] after all the data is copied. The installation server is now ready.

On Red Hat Enterprise Linux

This subsection provides the steps to prepare the NFS server on a system running Red Hat Enterprise Linux 5.2.

1. Move the DVD ISO image to an export directory:

```
mount -o loop /location/to/disk/RHEL5.2.iso /mnt/
cp -a /mnt/* /install/RHEL5.2/
```

- ```
umount /mnt/
```
2. Make sure the export directory is exported via NFS entry in /etc/exports. For example:  
`/install/RHEL5.2 *(ro, async, no_rootsquash)`
  3. Then restart the NFS daemon with:  
`/sbin/service nfs start`  
`/sbin/service nfs reload`

## Virtual optical device setup and installation

This installation option uses the virtual optical device on the Integrated Virtual Manager (IVM) to perform a CD/DVD installation of a Linux operating system image. The Linux image is stored in the IVM's virtual media library. Follow the steps given in 4.6.4, “Optical and Tape devices” on page 132 to copy the CD/DVD image onto the VIOS virtual media library. Once the CD/DVD image is copied onto the VIOS hard disk, assign the virtual optical device to the LPAR. The install process is the same as in 8.2, “Linux LPAR installation using DVD” on page 337 for Red Hat Enterprise Linux and 8.3, “Linux network installation (detailed)” on page 341 for SLES, but remember to select the *virtual optical device* as the boot device in the SMS menu and not the physical media tray.

**Note:** PowerVM must be installed and properly configured on the JS23 BladeCenter for this to work.

This installation is much faster than a traditional CD/DVD installation. The other benefit is that a read-only image in the VIOS media library is concurrently accessible to all the LPARs on the same IVM so you can run simultaneous installations of an operating system.



# Service and productivity tools for Linux

This appendix describes how to install IBM service diagnostic aids and productivity tools for the Linux operating system running on BladeCenter or IVM-managed servers for the JS23 BladeCenter.

This appendix contains the following:

- ▶ “IBM service diagnostic aids and productivity tools overview” on page 546
- ▶ “Install tools on Red Hat Enterprise Linux 5/SLES 11 running on BladeCenter servers” on page 548
- ▶ “Install tools on Red Hat Enterprise Linux 5/SLES 11 running on IVM-managed servers” on page 551

## **IBM service diagnostic aids and productivity tools overview**

The IBM service diagnostic and productivity packages for Linux on POWER architecture provide the latest system diagnostic information such as reliability, availability, and serviceability (RAS) functionality as well as the ability to modify logical partition (LPAR) profiles with hotplug, Dynamic Logical Partitioning (DLPAR), and Live Partition Migration capabilities.

The service and productivity packages available for your configuration are dependent on the system environment and the Linux operating system that is installed. The decision tree in Figure D-1 on page 547 shows how to determine the appropriate packages for your environment.

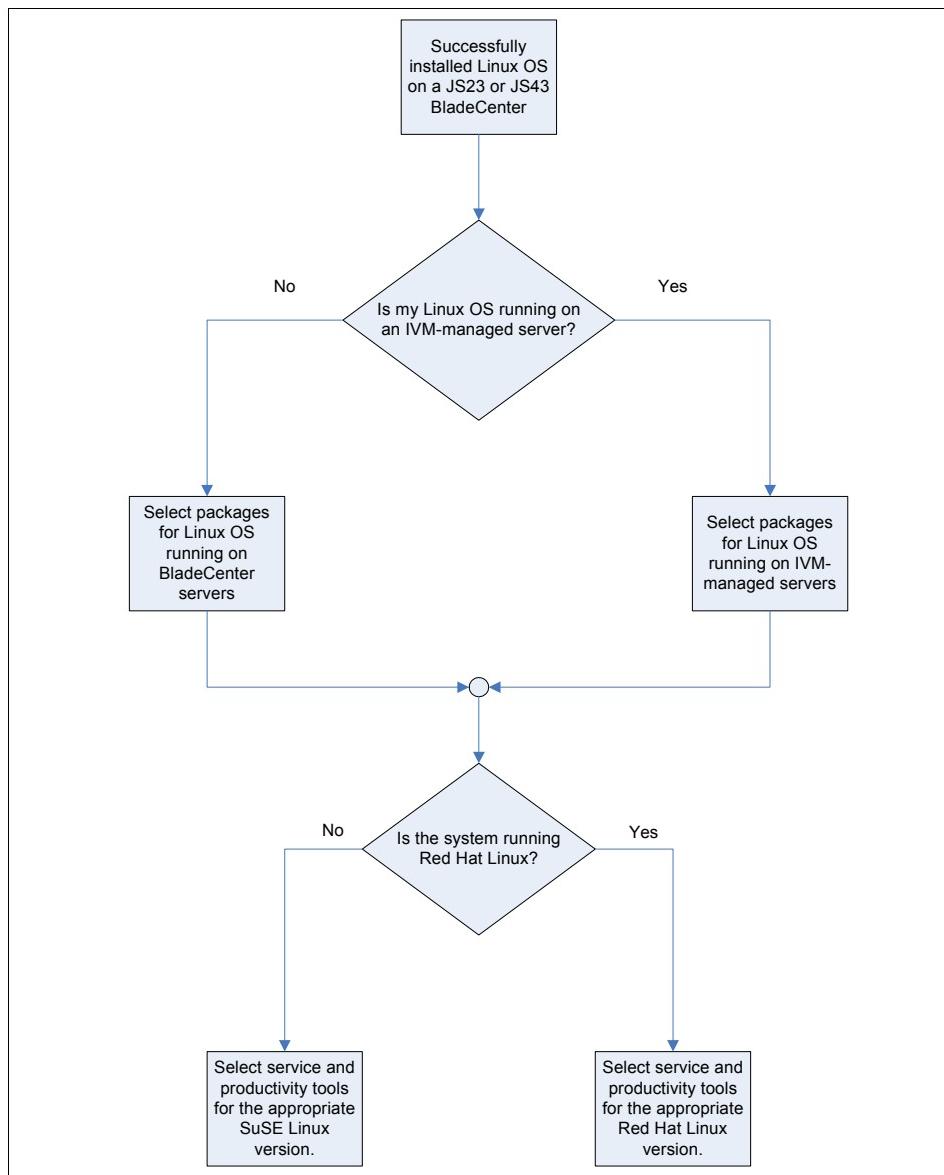


Figure D-1 Service diagnostic and productivity packages decision tree

# Install tools on Red Hat Enterprise Linux 5/SLES 11 running on BladeCenter servers

This section describes the steps to configure a JS23 BladeCenter running on a BladeCenter server with the service aids and productivity tools. These steps are applicable for systems running a native Red Hat Enterprise Linux 5/SLES 11 (or later) installation environment.

1. Use a Web browser to connect to

<https://www14.software.ibm.com/webapp/set2/sas/f/lopdiags/home.html>

2. On this Web site, click the **BladeCenter servers** link under the Red Hat or SuSE Linux distribution headings, as shown in Figure D-2.

Tools for Linux on POWER >

## Service and productivity tools

For Linux on POWER systems

Hardware service diagnostic aids and productivity tools, as well as installation aids for IBM servers running Linux operating systems on POWER6, POWER5 and POWER4 processors.

**Hardware diagnostic aids and productivity tools**

Tools and utilities are available for the following Linux distributions:

**Red Hat**

- on BladeCenter servers
- on HMC- or IVM-managed servers
- on other servers

**Suse Linux**

- on BladeCenter servers
- on HMC- or IVM-managed servers
- on other servers

Figure D-2 Select “on BladeCenter servers”

3. Click the tab (as shown in Figure D-3) that matches your Linux operating system (OS) level. The packages under each tab are unique to that Linux OS level.

Tools for Linux on POWER >

## Service and productivity tools

For Red Hat Linux on Blade servers

RHEL 5    RHEL 4    RHEL 3

The following tools are available for BladeCenter servers running Red Hat Linux RHEL 5. The Tool name link provides a brief description of the tool.

Tool packages must be installed in the order listed in the table.

*Figure D-3 OS level selection tabs example*

4. Click and save each of the packages under the Package downloads column.  
At the time of this publication the packages were:

| RHEL 5 tools                |                                                                                                                                                                                         |             |
|-----------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------|
| Tool name                   | Package downloads                                                                                                                                                                       | Last update |
| Platform Enablement Library | libertas-1.3.4-0.ppc64.rpm                                                                                                                                                              | 2009.03.05  |
| Hardware Inventory          | lsvpd-1.6.5-1.ppc64.rpm<br>libvpd-2.1.0-1.ppc64.rpm<br>(prerequisite for <b>lsvpd</b> RPM)                                                                                              | 2009.03.05  |
| Service log                 | servicelog-1.0.1-1.ppc64.rpm<br>libservicelog-1.0.1-1.ppc64.rpm<br>(prereq for <b>servicelog</b> RPM)<br>libservicelog-devel-1.0.1-1.ppc64.rpm<br>(prereq for <b>libservicelog</b> RPM) | 2009.03.10  |
| Error Log Analysis          | diagela-2.2.3-0.ppc64.rpm                                                                                                                                                               | 2009.02.24  |

*Figure D-4 Available packages for Red Hat on BladeCenter servers*

| Tool                        | Basic Information                                                                                 |
|-----------------------------|---------------------------------------------------------------------------------------------------|
| Platform Enablement Library | A library that allows applications to access certain functionality provided by platform firmware. |
| Hardware Inventory          | Provides Vital Product Data (VPD) about hardware components to higher-level serviceability tools. |

| Tool               | Basic Information                                                                         |
|--------------------|-------------------------------------------------------------------------------------------|
| Service Log        | Creates a database to store system-generated events that may require service.             |
| Error Log Analysis | Provides automatic analysis and notification of errors reported by the platform firmware. |

**Tip:** Click the links under the Tool name column for the latest detailed description of each tool.

5. Use a transfer protocol such as FTP or SCP to send each \*.rpm package to the target system or save these rpm packages to a CD or DVD and mount the device (see the CD/DVD tip below).
6. Install each rpm package with  

```
rpm -Uvh <packagename>.rpm
```

**Important:** These packages *must* be installed in the order listed in the table. Otherwise there will be dependency failures.
7. We recommend a system shutdown and restart after installing all the packages.
8. The service aids and productivity tools section is complete.

**Tip1:** If you are placing the rpms on a CD and DVD, here are some steps to access the files:

1. Assign the JS23 media tray to the appropriate blade bay.
2. Mount the media tray to a directory on the system such as /mnt on Red Hat:  

```
mount /dev/cdrom /mnt/
```

 or on SuSE Linux:  

```
mount /dev/sr0 /mnt/
```
3. Move to the mounted directory with  

```
cd /mnt
```
4. Install each rpm with  

```
rpm -Uvh <packagename>.rpm
```

**Tip2:** We recommend placing these rpms in a yum repository to quickly update or install these tools on a large number of machines.

## Install tools on Red Hat Enterprise Linux 5/SLES 11 running on IVM-managed servers

This section describes the steps to configure a JS23 BladeCetner LPAR running on a IVM-managed server with the service aids and productivity tools.

1. Use a Web browser to connect to

<https://www14.software.ibm.com/webapp/set2/sas/f/lopdiags/home.html>

2. On this Web site, click **HMC- or IVM-managed servers** in the Red Hat/SuSE Linux section, depending on what OS is running on the LPAR.
3. Click the tab (as shown in Figure D-3 on page 549) that matches your Linux operating system (OS) level. The packages under each tab are unique to that Linux OS level.
4. Click on and save each of the packages under the Download column. At the time of this publication, the packages are as shown in Figure D-5 on page 552.

| <b>RHEL 5 tools</b>           |                                               |
|-------------------------------|-----------------------------------------------|
| <b>Tool name</b>              | <b>Download</b>                               |
| Platform Enablement Library   | libertas-1.3.3-0.ppc64.rpm                    |
| SRC                           | src-1.3.0.2-07305.ppc.rpm                     |
| RSCT utilities                | rsct.core.utils-2.5.0.1-07305.ppc.rpm         |
| RSCT core                     | rsct.core-2.5.0.1-07305.ppc.rpm               |
| CSM core                      | csm.core-1.7.0.1-57.ppc.rpm                   |
| CSM client                    | csm.client-1.7.0.1-57.ppc.rpm                 |
| ServiceRM                     | devices.chrp.base.ServiceRM-2.2.0.0-6.ppc.rpm |
| DynamicRM                     | DynamicRM-1.3.2-0.ppc64.rpm                   |
| Hardware Inventory            | lsvpd-0.15.1-1.ppc.rpm                        |
| Service log                   | servicelog-0.2.9-0.ppc64.rpm                  |
| Error Log Analysis            | diagela-2.2.2-0.ppc64.rpm                     |
| PCI Hotplug Tools             | rpa-pci-hotplug-1.0-30.ppc64.rpm              |
| Dynamic Reconfiguration Tools | rpa-dlpar-1.0-47.ppc64.rpm                    |
| Inventory Scout               | IBMinvscout-2.2-5.ppc.rpm                     |

Figure D-5 Available packages for Red Hat/SuSE Linux on IVM-managed server

| Tool                        | Basic Information                                                                                                                                                                                                                                    |
|-----------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Platform Enablement Library | A library that allows application to access certain functionality provided by platform firmware.                                                                                                                                                     |
| SRC                         | Manages daemons on the systems.                                                                                                                                                                                                                      |
| RSCT utilities              | RSC packages provide the Resource Monitoring and Control (RMC) functions and infrastructure needed to monitor and manage one or more Linux systems                                                                                                   |
| RSCT core                   | See description above.                                                                                                                                                                                                                               |
| CSM core                    | CSM packages provide for the exchange of host-based authentication security keys.                                                                                                                                                                    |
| CSM-client                  | See description above.                                                                                                                                                                                                                               |
| ServiceRM                   | Service Resource Manager is a Reliable, Scalable, Cluster Technology (RSCT) resource manager that creates the Serviceable Events from the output of the Error Log Analysis Tool (diagela).                                                           |
| DynamicRM                   | Dynamic Resource Manager is a Reliable, Scalable, Cluster Technology (RSCT) resource manager that allows a IVM to dynamically add or remove processors or I/O slots from a running partition and perform certain shutdown operations on a partition. |
| Hardware Inventory          | Provides Vital Product Data (VPD) about hardware components to higher-level serviceability tools.                                                                                                                                                    |
| Service Log                 | Creates a database to store system-generated events that may require service.                                                                                                                                                                        |

| Tool                         | Basic Information                                                                         |
|------------------------------|-------------------------------------------------------------------------------------------|
| Error Log Analysis           | Provides automatic analysis and notification of errors reported by the platform firmware. |
| PCI Hotplug Tools            | Allows PCI devices to be added, removed, or replaced while the system is in operation.    |
| Dynamic Reconfiguration Tool | Allows the addition and removal of processors and I/O slots from a running partition.     |
| Inventory Scout              | Surveys one or more systems for hardware and software information.                        |

**Tip:** Click the links under the Tool name column for the latest detailed description of each tool.

5. Use a transfer protocol such as FTP or SCP to send each \*.rpm package to the target system or save them to a CD or DVD and mount the device (see “Tip1: If you are placing the rpms on a CD and DVD, here are some steps to access the files.” on page 550).

6. Install each rpm package with

```
rpm -Uvh <packagename>.rpm
```

**Important:** These packages *must* be installed in the order listed in the table. Otherwise there will be dependency failures.

7. We recommend a complete LPAR shutdown and reactivation after installing all the packages.
8. Log on to your IVM and ensure that the General tab of the LPAR properties shows all DLPAR services enabled with Yes, as Figure D-6 on page 555 shows.

| Dynamic Logical Partitioning (DLPAR) |              |
|--------------------------------------|--------------|
| Partition hostname or IP address:    | 172.16.1.196 |
| Partition communication state:       | Active       |
| Memory DLPAR capable:                | Yes          |
| Processing DLPAR capable:            | Yes          |
| I/O adapter DLPAR capable:           | Yes          |

*Figure D-6 DLPAR and Live Partition mobility services are enabled*

See Chapter 4, “System planning and configuration using VIOS with IVM” on page 71 for more information on IVM options and functions.

9. Installation of the service aids and productivity tools is complete.

**Tip:** We recommend placing these rpms in a yum repository to quickly update or install these tools on a large number of machines.



# Abbreviations and acronyms

|             |                                         |               |                                       |
|-------------|-----------------------------------------|---------------|---------------------------------------|
| <b>ABR</b>  | Automatic BIOS recovery                 | <b>CCSP</b>   | Cisco Certified Security Professional |
| <b>AC</b>   | alternating current                     | <b>CD-ROM</b> | compact disc read only memory         |
| <b>ACL</b>  | access control list                     | <b>CDP</b>    | Cisco Discovery Protocol              |
| <b>AES</b>  | Advanced Encryption Standard            | <b>CE</b>     | Conformité Européene                  |
| <b>AMD™</b> | Advanced Micro Devices™                 | <b>CLI</b>    | command-line interface                |
| <b>AMM</b>  | Advanced Management Module              | <b>CNA</b>    | Cisco Network Assistance              |
| <b>API</b>  | application programming interface       | <b>CNS</b>    | Cisco Network Services                |
| <b>APV</b>  | Advanced Power Virtualization           | <b>COG</b>    | configuration and option guide        |
| <b>ARP</b>  | Address Resolution Protocol             | <b>CPM</b>    | Copper Pass-thru Module               |
| <b>AS</b>   | Australian Standards                    | <b>CPU</b>    | central processing unit               |
| <b>ASF</b>  | Alert Standard Format                   | <b>CRU</b>    | customer replaceable units            |
| <b>ASIC</b> | application-specific integrated circuit | <b>CSM</b>    | Cluster Systems Management            |
| <b>ASR</b>  | automatic server restart                | <b>CTS</b>    | clear to send                         |
| <b>BASP</b> | Broadcom Advanced Server Program        | <b>DASD</b>   | direct access storage device          |
| <b>BBI</b>  | browser-based interface                 | <b>DC</b>     | domain controller                     |
| <b>BCM</b>  | Broadcom                                | <b>DDM</b>    | Deployment and Management             |
| <b>BE</b>   | Broadband Engine                        | <b>DDR</b>    | Double Data Rate                      |
| <b>BGP</b>  | Border Gateway Protocol                 | <b>DHCP</b>   | Dynamic Host Configuration Protocol   |
| <b>BIOS</b> | basic input output system               | <b>DIMM</b>   | dual inline memory module             |
| <b>BMC</b>  | Baseboard Management Controller         | <b>DIP</b>    | destination IP                        |
| <b>BNT</b>  | BLADE Network Technologies, Inc         | <b>DMAC</b>   | destination MAC address               |
| <b>BOFM</b> | BladeCenter Open Fabric Manager         | <b>DNS</b>    | Domain Name System                    |
| <b>BPDU</b> | Bridge protocol data unit               | <b>DP</b>     | dual processor                        |
| <b>BSE</b>  | BladeCenter Storage Expansion           | <b>DPOD</b>   | Dynamic Ports on Demand               |
| <b>BSMP</b> | blade system management processor       | <b>DSA</b>    | Dynamic System Analysis               |
| <b>BTU</b>  | British Thermal Unit                    | <b>DSCP</b>   | Differentiated Services Code Point    |
| <b>CCDA</b> | Cisco Certified Design Associate        | <b>DSM</b>    | disk storage module                   |
| <b>CCNP</b> | Cisco Certified Network Professional    | <b>DSUB</b>   | D-subminiature                        |
|             |                                         | <b>DTP</b>    | Dynamic Trunking Protocol             |
|             |                                         | <b>DVI</b>    | Digital Video Interface               |

|                |                                               |               |                                                   |
|----------------|-----------------------------------------------|---------------|---------------------------------------------------|
| <b>DVMRP</b>   | Distance Vector Multicast Routing Protocol    | <b>HSDC</b>   | high speed daughter card                          |
| <b>DVS</b>     | Digital Video Surveillance                    | <b>HSFF</b>   | high speed Ethernet switch module                 |
| <b>ECC</b>     | error checking and correcting                 | <b>HSIBPM</b> | high-speed form factor                            |
| <b>EDA</b>     | Electronic Design Automation                  |               | high-speed InfiniBand pass-thru module            |
| <b>EIGRP</b>   | Enhanced Interior Gateway Routing Protocol    | <b>HSIBSM</b> | high speed InfiniBand switch module               |
| <b>EMC</b>     | electromagnetic compatibility                 | <b>HSRP</b>   | Hot Standby Routing Protocol                      |
| <b>EMEA</b>    | Europe, Middle East, Africa                   | <b>HT</b>     | Hyper-Threading                                   |
| <b>EOT</b>     | Enhanced object tracking                      | <b>HTTP</b>   | Hypertext Transfer Protocol                       |
| <b>EPOW</b>    | Early Power Off Warning                       | <b>I/O</b>    | input/output                                      |
| <b>ESD</b>     | electrostatic discharge                       | <b>IB</b>     | InfiniBand                                        |
| <b>ESM</b>     | Ethernet switch modules                       | <b>IBBM</b>   | InfiniBand bridge module                          |
| <b>ETSI</b>    | European Telecommunications Standard Industry | <b>IBM</b>    | International Business Machines                   |
| <b>FAN</b>     | Fabric Address Notification                   | <b>ICMP</b>   | Internet control message protocol                 |
| <b>FB-DIMM</b> | Fully Buffered DIMMs                          | <b>ICPM</b>   | Intelligent Copper Pass-thru Module               |
| <b>FBDIMM</b>  | Fully Buffered DIMM                           | <b>ID</b>     | identifier                                        |
| <b>FC</b>      | Fibre Channel                                 | <b>IDE</b>    | integrated drive electronics                      |
| <b>FCP</b>     | Flow Control Packet                           | <b>IEC</b>    | International Electro-technical Commission        |
| <b>FCSM</b>    | Fibre Channel Switch Module                   |               |                                                   |
| <b>FDD</b>     | floppy diskette drive                         | <b>IEEE</b>   | Institute of Electrical and Electronics Engineers |
| <b>FDX</b>     | full duplex                                   |               |                                                   |
| <b>FSB</b>     | front-side bus                                | <b>IGESM</b>  | Intelligent Gigabit Ethernet Switch Module        |
| <b>FTP</b>     | File Transfer Protocol                        | <b>IGMP</b>   | Internet Group Management Protocol                |
| <b>FTSS</b>    | Field Technical Sales Support                 |               |                                                   |
| <b>GB</b>      | gigabyte                                      | <b>IGRP</b>   | Interior Gateway Routing Protocol                 |
| <b>GUI</b>     | graphical user interface                      | <b>IM</b>     | instant messaging                                 |
| <b>HA</b>      | high availability                             | <b>IME</b>    | Integrated Mirroring Enhanced                     |
| <b>HBA</b>     | host bus adapter                              | <b>IOS</b>    | Internetwork Operating System                     |
| <b>HCA</b>     | host channel adapter                          | <b>IP</b>     | Internet Protocol                                 |
| <b>HD</b>      | high definition                               | <b>IPM</b>    | Intelligent Pass-thru Module                      |
| <b>HDD</b>     | hard disk drive                               | <b>IPMI</b>   | Intelligent Platform Management Interface         |
| <b>HH</b>      | half high                                     |               |                                                   |
| <b>HPC</b>     | high performance computing                    | <b>IPTV</b>   | Internet Protocol Television                      |
| <b>HS</b>      | hot swap                                      | <b>IRDP</b>   | ICMP Router Discovery Protocol                    |
|                |                                               | <b>IS</b>     | information store                                 |

|             |                                              |               |                                            |
|-------------|----------------------------------------------|---------------|--------------------------------------------|
| <b>ISL</b>  | Inter-Switch Link                            | <b>MVR</b>    | Multicast VLAN registration                |
| <b>ISMP</b> | Integrated System Management Processor       | <b>NAT</b>    | Network Address Translation                |
| <b>ISP</b>  | Internet service provider                    | <b>NDCLA</b>  | Non-Disruptive Code Load Activation        |
| <b>IT</b>   | information technology                       | <b>NEBS</b>   | Network Equipment Building System          |
| <b>ITS</b>  | IBM Integrated Technology Services           | <b>NGN</b>    | next-generation network                    |
| <b>ITSO</b> | International Technical Support Organization | <b>NIC</b>    | network interface card                     |
| <b>IVM</b>  | Integrated Virtualization Manager            | <b>NMI</b>    | non-maskable interrupt                     |
| <b>KB</b>   | kilobyte                                     | <b>NOS</b>    | network operating system                   |
| <b>KVM</b>  | keyboard video mouse                         | <b>NPIV</b>   | N_Port ID Virtualization                   |
| <b>LACP</b> | Link Aggregation Control Protocol            | <b>NSF</b>    | Notes Storage File                         |
| <b>LAN</b>  | local area network                           | <b>NTP</b>    | Network Time Protocol                      |
| <b>LED</b>  | light emitting diode                         | <b>OBFL</b>   | On-board failure logging                   |
| <b>LLDP</b> | Link Layer Discovery Protocol                | <b>ODPA</b>   | On-Demand Port Activation                  |
| <b>LPAR</b> | logical partitions                           | <b>OFED</b>   | OpenFabrics Enterprise Distribution        |
| <b>LPH</b>  | low profile handle                           | <b>OFM</b>    | Open Fabric Manager                        |
| <b>LR</b>   | long range                                   | <b>OPM</b>    | Optical Pass-thru Module                   |
| <b>LTO</b>  | Linear Tape-Open                             | <b>OS</b>     | operating system                           |
| <b>LUN</b>  | logical unit number                          | <b>OSPF</b>   | Open Shortest Path First                   |
| <b>MAC</b>  | media access control                         | <b>PBR</b>    | Policy-based routing                       |
| <b>MAN</b>  | metropolitan area network                    | <b>PC</b>     | personal computer                          |
| <b>MB</b>   | megabyte                                     | <b>PCI</b>    | Peripheral Component Interconnect          |
| <b>MDS</b>  | Multilayer DataCenter Switch                 | <b>PDF</b>    | Portable Document Format                   |
| <b>MIB</b>  | management information base                  | <b>PDU</b>    | power distribution unit                    |
| <b>MIO</b>  | Memory and I/O                               | <b>PFA</b>    | Predictive Failure Analysis                |
| <b>MM</b>   | Management Module                            | <b>POST</b>   | power-on self test                         |
| <b>MMF</b>  | Multi Mode Fiber                             | <b>PPP</b>    | point-to-point protocol                    |
| <b>MP</b>   | multiprocessor                               | <b>PVST</b>   | Per-VLAN Spanning Tree                     |
| <b>MPE</b>  | Multi Processor Expansion                    | <b>PXE</b>    | Preboot Execution Environment              |
| <b>MPI</b>  | Message Passing Interface                    | <b>RAC</b>    | Real Application Clusters                  |
| <b>MSDP</b> | Multicast Source Discovery Protocol          | <b>RADIUS</b> | Remote Authentication Dial In User Service |
| <b>MSIM</b> | Multi-Switch Interconnect Module             | <b>RAID</b>   | redundant array of independent disks       |
| <b>MSTP</b> | Multiple Spanning Tree Protocol              |               |                                            |
| <b>MTM</b>  | machine-type-model                           | <b>RAM</b>    | random access memory                       |

|                |                                                            |               |                                                  |
|----------------|------------------------------------------------------------|---------------|--------------------------------------------------|
| <b>RAS</b>     | remote access services; row address strobe                 | <b>SIP</b>    | source IP                                        |
| <b>RDAC</b>    | Redundant Disk Array Controller                            | <b>SLB</b>    | Server Load Balancing                            |
| <b>RDC</b>     | Remote Desktop Connection                                  | <b>SLES</b>   | SUSE Linux Enterprise Server                     |
| <b>RDIMM</b>   | registered DIMM                                            | <b>SMAC</b>   | source MAC address                               |
| <b>RDM</b>     | Remote Deployment Manager                                  | <b>SMI-S</b>  | Storage Management Initiative - Specification    |
| <b>RDMA</b>    | Remote Direct Memory Access                                | <b>SMP</b>    | symmetric multiprocessing                        |
| <b>RETAIN®</b> | Remote Electronic Technical Assistance Information Network | <b>SMS</b>    | System Management Services                       |
| <b>RHEL</b>    | Red Hat Enterprise Linux                                   | <b>SNMP</b>   | Simple Network Management Protocol               |
| <b>RIP</b>     | Routing Information Protocol                               | <b>SOL</b>    | Serial over LAN                                  |
| <b>RMCP</b>    | Remote Management Control Protocol                         | <b>SPORE</b>  | ServerProven Opportunity Request for Evaluation  |
| <b>RMON</b>    | Remote Monitoring                                          | <b>SR</b>     | short range                                      |
| <b>RP</b>      | route processor                                            | <b>SRP</b>    | Storage RDMA Protocol                            |
| <b>RPF</b>     | reverse path forwarding                                    | <b>SRR</b>    | shaped round robin                               |
| <b>RPM</b>     | revolutions per minute                                     | <b>SSCT</b>   | Standalone Solution Configuration Tool           |
| <b>RPQ</b>     | Request Per Qualification                                  |               |                                                  |
| <b>RSA</b>     | Remote Supervisor Adapter                                  | <b>SSD</b>    | solid state drive                                |
| <b>RSCN</b>    | Registered State Change Notification                       | <b>SSH</b>    | Secure Shell                                     |
|                |                                                            | <b>SSL</b>    | Secure Sockets Layer                             |
| <b>RSTP</b>    | Rapid Spanning Tree Protocol                               | <b>SSP</b>    | Serial SCSI Protocol                             |
| <b>SAN</b>     | storage area network                                       | <b>STP</b>    | Spanning Tree Protocol                           |
| <b>SAS</b>     | Serial Attached SCSI                                       | <b>TACACS</b> | Terminal Access Controller Access Control System |
| <b>SASCM</b>   | SAS Connectivity Module                                    |               |                                                  |
| <b>SATA</b>    | Serial ATA                                                 | <b>TB</b>     | terabyte                                         |
| <b>SBB</b>     | Sales Building Block                                       | <b>TCO</b>    | total cost of ownership                          |
| <b>SCM</b>     | Supply Chain Management                                    | <b>TCP</b>    | Transmission Control Protocol                    |
| <b>SCSI</b>    | Small Computer System Interface                            | <b>TCP/IP</b> | Transmission Control Protocol/Internet Protocol  |
| <b>SDD</b>     | Subsystem Device Driver                                    | <b>TFTP</b>   | Trivial File Transfer Protocol                   |
| <b>SDK</b>     | Software Developers' Kit                                   | <b>TSM</b>    | Tivoli Storage Manager                           |
| <b>SDR</b>     | Single Data Rate                                           | <b>TX</b>     | transmit                                         |
| <b>SDRAM</b>   | static dynamic RAM                                         | <b>UDLD</b>   | UniDirectional link detection                    |
| <b>SFF</b>     | Small Form Factor                                          | <b>UDP</b>    | user datagram protocol                           |
| <b>SFP</b>     | small form-factor pluggable                                | <b>ULP</b>    | upper layer protocols                            |
| <b>SIMD</b>    | single instruction multiple data                           | <b>URL</b>    | Uniform Resource Locator                         |
| <b>SIO</b>     | Storage and I/O                                            |               |                                                  |

|             |                                    |
|-------------|------------------------------------|
| <b>USB</b>  | universal serial bus               |
| <b>UTF</b>  | Universal Telco Frame              |
| <b>UTP</b>  | unshielded twisted pair            |
| <b>VBS</b>  | Virtual Blade Switch               |
| <b>VGA</b>  | video graphics array               |
| <b>VIOS</b> | Virtual I/O Server                 |
| <b>VLAN</b> | virtual LAN                        |
| <b>VLP</b>  | very low profile                   |
| <b>VM</b>   | virtual machine                    |
| <b>VMPS</b> | VLAN Membership Policy Server      |
| <b>VNC</b>  | Virtual Network Computing          |
| <b>VOIP</b> | Voice over Internet Protocol       |
| <b>VPD</b>  | vital product data                 |
| <b>VPN</b>  | virtual private network            |
| <b>VQP</b>  | VLAN Query Protocol                |
| <b>VRRP</b> | virtual router redundancy protocol |
| <b>VSAN</b> | Virtual Storage Area Network       |
| <b>VT</b>   | Virtualization Technology          |
| <b>VTP</b>  | VLAN Trunking Protocol             |
| <b>WAN</b>  | wide area network                  |
| <b>WOL</b>  | Wake on LAN®                       |
| <b>WTD</b>  | Weighted tail drop                 |
| <b>WWN</b>  | World Wide Name                    |
| <b>XDR</b>  | extreme data rate                  |
| <b>XM</b>   | extended memory                    |



# Related publications

The publications listed in this section are considered particularly suitable for a more detailed discussion of the topics covered in this book.

## IBM Redbooks

For information about ordering these publications, see “How to get Redbooks” on page 568. Note that some of the documents referenced here may be available in softcopy only.

- ▶ *IBM BladeCenter Products and Technology*, SG24-7523
- ▶ *IBM System i and System p*, SG24-7487
- ▶ *IBM System Storage DS4000 and Storage Manager V10.10*, SG24-7010
- ▶ *PowerVM Live Partition Mobility on IBM System p*, SG24-7460
- ▶ *Managing OS/400 with Operations Navigator V5R1 Volume 1: Overview and More*, SG24-6226
- ▶ *Managing OS/400 with Operations Navigator V5R1 Volume 5: Performance Management*, SG24-6565
- ▶ *PowerVM Virtualization on IBM System p Managing and Monitoring*, SG24-7590
- ▶ *Implementing IBM Director 5.20*, SG24-6188
- ▶ *Going Green with IBM Active Energy Manager*, REDP-4361
- ▶ *Integrated Virtualization Manager on IBM System p5*, REDP-4061
- ▶ *Implementing the IBM BladeCenter S Chassis*, SG24-76827
- ▶ *IBM System i Overview: Models 515, 525, 550, 570, 595, and More*, REDP-5052
- ▶ *IBM BladeCenter JS12 and JS22 Implementation Guide*, SG247655

## Online resources

These Web sites are also relevant as further information sources:

- ▶ IBM Systems Director Active Energy Manager Version 3.1.1 is an IBM Director extension. For more information about the IBM Active Energy Manager see:  
<http://www.ibm.com/systems/management/director/extensions/actengmrg.html>
- ▶ IBM periodically releases maintenance packages for the AIX 5L operating system. These packages are available on CD-ROM, or you can download them from the following Web site:  
<http://www.ibm.com/eserver/support/fixes/fixcentral/main/pseries/aix>
- ▶ In AIX 5L V5.3, the **suma** command is also available, which helps the administrator to automate the task of checking and downloading operating system downloads. For more information about the suma command functionality, visit the following Web site:  
<http://www.ibm.com/systems/p/os/aix/whitepapers/suma.html>
- ▶ For information about the features and external devices supported by Linux on the IBM BladeCenter JS23 and IBM BladeCenter JS43 servers visit the following Web site:  
<http://www.ibm.com/systems/power/software/linux/index.html>
- ▶ For information about SUSE Linux Enterprise Server 11, visit the following Web site:  
<http://developer.novell.com/yesssearch/Search.jsp>
- ▶ For information about Red Hat Enterprise Linux, visit the following Web site:  
<https://hardware.redhat.com/?pagename=hcl&view=certified&vendor=4&class=8>
- ▶ Many of the features described in this document are operating system dependent and might not be available on Linux. For more information, visit the following Web site:  
[http://www.ibm.com/systems/p/software/whitepapers/linux\\_overview.html](http://www.ibm.com/systems/p/software/whitepapers/linux_overview.html)
- ▶ For more information of IBM i 6.1 operating systems running on IBM BladeCenter JS23 and IBM BladeCenter JS43 see:  
<http://www.ibm.com/systems/power/hardware/blades/ibmi.html>
- ▶ Visit the following Web site to download IBM Director:  
<http://www-03.ibm.com/systems/management/director/downloads/>
- ▶ Complete VIOS configuration and maintenance information can be found in *System i and System p Using the Virtual I/O Server*, found at this link:

[http://publib.boulder.ibm.com/infocenter/iseries/v1r3s/en\\_US/info/iphb1/iphb1.pdf](http://publib.boulder.ibm.com/infocenter/iseries/v1r3s/en_US/info/iphb1/iphb1.pdf)

- ▶ The BladeCenter Interoperability Guide can be found at:  
<https://www-304.ibm.com/systems/support/supportsite.wss/docdisplay?lnodocid=MIGR-5073016&brandind=50000020>
  - ▶ The Virtual I/O server data sheet gives an overview of supported storage subsystems and the failover driver that is supported with the subsystem. The data sheet can be found at:  
<http://www14.software.ibm.com/webapp/set2/sas/f/vios/documentation/datasheet.html>
  - ▶ All supported hardware and operating systems are listed on IBM ServerProven. ServerProven can be found at:  
<http://www-03.ibm.com/servers/eserver/serverproven/compat/us/>
  - ▶ Verify the supported operating systems on the blade by using the link to NOS Support on the Compatibility for BladeCenter products page of ServerProven at:  
<http://www-03.ibm.com/servers/eserver/serverproven/compat/us/eserver.html>
  - ▶ To verify which of those components supported by the blade are supported by the Virtual IO server as well. The data sheet can be found at:  
<http://www14.software.ibm.com/webapp/set2/sas/f/vios/documentation/datasheet.html>
  - ▶ The DS8000 interoperability matrix can be found at:  
<http://www.ibm.com/servers/storage/disk/ds8000/interop.html>
  - ▶ The DS6000 interoperability matrix can be found at:  
<http://www.ibm.com/servers/storage/disk/ds6000/interop.html>
  - ▶ The DS4000 interoperability matrix can be found at:  
<http://www.ibm.com/servers/storage/disk/ds4000/interop-matrix.html>
  - ▶ The DS3000 interoperability matrix can be found at:  
<http://www-03.ibm.com/systems/storage/disk/ds3000/pdf/interop.pdf>
  - ▶ The ESS interoperability matrix can be found at:  
<http://www.ibm.com/servers/storage/disk/ess/interop-matrix.html>
- The N series interoperability matrix can be found at:  
<http://www-03.ibm.com/systems/storage/nas/interophome.html>
- ▶ The SAN volume Controller support matrix can be found at:

<http://www-304.ibm.com/jct01004c/systems/support/supportsite.wss/supportresources?taskkind=3&brandind=5000033&familyind=5329743>

- ▶ The SAN switch interoperability matrix can be found at:  
<http://www-03.ibm.com/systems/storage/san/index.html>
- ▶ The System Storage Interoperation Center (SSIC) helps to identify a supported storage environment. You find this web based tool at:  
<http://www-03.ibm.com/systems/support/storage/config/ssic>
- ▶ DS3000 and DS4000 support Web pages usually provide update packages for the supported adapters that contain the settings required for the HBA. They can be found at:  
<http://www-03.ibm.com/systems/support/storage/config/hba/index.wss>
- ▶ The Storage Configuration Manager (SCM) is proposed to create an individual configuration if you are not familiar using the SAS I/O module command line interface. The SCM software can be downloaded from:  
<http://www-947.ibm.com/systems/support/supportsite.wss/selectproduct?familyind=5370552&osind=0&brandind=5000016&oldbrand=5000016&oldfamily=5370552&oldtype=0&taskkind=2&psid=sp&continue.x=15&continue.y=15>
- ▶ For more information to IBM System i Access for Windows V6R1 see:  
<http://www.ibm.com/systems/i/software/access/index.html>
- ▶ To obtain the IBM System i Access for Windows software:  
<http://www.ibm.com/systems/i/software/access/caorder.html>
- ▶ Service tools user IDs are created through DST or SST and are separate from IBM i 6.1 user profiles. Refer to Service tools user IDs at:  
<http://publib.boulder.ibm.com/infocenter/systems/scope/i5os/topic/rzamh/rzamhwwhatuserids.htm>
- ▶ The IBM Workload Estimator is available at:  
<http://www.ibm.com/systems/support/tools/estimator/index.html>
- ▶ Performance actions related to disk formatting are described in InfoCenter at the following:  
<https://publib.boulder.ibm.com/infocenter/systems/scope/i5os/index.jsp?topic=/rzajy/rzajoverview.html>
- ▶ Language feature codes are available at:  
<http://publib.boulder.ibm.com/infocenter/scope/i5os/topic/rzahc/rzahnlvfeaturecodes.htm#rzhcnlvfeaturecodes>
- ▶ Software installation information for IBM i:

<http://publib.boulder.ibm.com/infocenter/iseries/v5r3/topic/rzahc/rzahcsinstallprocess.htm>

- ▶ Use the IBM i recommended Fixes Web site to get a list of the latest recommended PTFs:

[http://www-912.ibm.com/s\\_dir/slkbbase.nsf/recommendedfixes](http://www-912.ibm.com/s_dir/slkbbase.nsf/recommendedfixes)

- ▶ The primary Web site for downloading fixes for all operating systems and applications refer to:

<http://www-912.ibm.com/eserver/support/fixes>

- ▶ More detailed information to the IBM Systems Director Navigator for i functionality can be found at:

<http://www.ibm.com/systems/i/software/navigator/index.html>

- ▶ A good source for processor and memory requirements for PowerVM partitions based on I/O requirements is found at:

<http://www14.software.ibm.com/webapp/set2/sas/f/vios/documentation/perf.html>

- ▶ The Red Hat Recommended Partitioning Scheme is available at:

[http://www.redhat.com/docs/en-US/Red\\_Hat\\_Enterprise\\_Linux/5.2/html/Installation\\_Guide/ch11s03.html](http://www.redhat.com/docs/en-US/Red_Hat_Enterprise_Linux/5.2/html/Installation_Guide/ch11s03.html)

- ▶ A source for processor and memory requirements for PowerVM partitions based on I/O requirements is found at:

<http://www14.software.ibm.com/webapp/set2/sas/f/vios/documentation/perf.html>

- ▶ The Red Hat Recommended Partitioning Scheme is available at:

[http://www.redhat.com/docs/en-US/Red\\_Hat\\_Enterprise\\_Linux/5.2/html/Installation\\_Guide/ch11s03.html](http://www.redhat.com/docs/en-US/Red_Hat_Enterprise_Linux/5.2/html/Installation_Guide/ch11s03.html)

- ▶ RedHat installation instructions are available here:

[http://www.redhat.com/docs/en-US/Red\\_Hat\\_Enterprise\\_Linux/5.2/html/Installation\\_Guide/pt-install-info-ppc.html](http://www.redhat.com/docs/en-US/Red_Hat_Enterprise_Linux/5.2/html/Installation_Guide/pt-install-info-ppc.html)

- ▶ A good source for processor and memory requirements for PowerVM partitions based on I/O requirements is found at:

<http://www14.software.ibm.com/webapp/set2/sas/f/vios/documentation/perf.html>

- ▶ The Novell Web site has additional installation preparation information for SLES10 Sp2 available at:

<http://www.novell.com/documentation/sles10/index.html>

- ▶ Your awesome Setup Tool (YaST) will assist with the completion of a SLES install. More detailed installation instructions are available here:  
[http://www.novell.com/documentation/sles10/sles\\_admin/index.html?page=/documentation/sles10/sles\\_admin/data/sles\\_admin.html](http://www.novell.com/documentation/sles10/sles_admin/index.html?page=/documentation/sles10/sles_admin/data/sles_admin.html)
- ▶ The link to the Virtual I/O Server download site is also available here:  
<http://techsupport.services.ibm.com/server/vios/download>
- ▶ Linux partitions must have the Dynamic Reconfiguration Tools package for HMC or IVM managed servers installed from the Service and productivity tools Web site available at:  
<https://www14.software.ibm.com/webapp/set2/sas/f/lopdiags/home.html>
- ▶ For details about how to setup SOL see the Serial over LAN setup Guide. This guide can be found at:  
<http://www-304.ibm.com/systems/support/supportsite.wss/docdisplay?ln docid=MIGR-54666&brandind=5000020>
- ▶ The *Management Module command-line interface reference guide* can you find online at:  
<http://www-304.ibm.com/systems/support/supportsite.wss/docdisplay?ln docid=MIGR-54667&brandind=5000020>
- ▶ For a technical overview and full details, as well as latest updates on IBM i on Power blades, see the Read-me First:  
<http://www-03.ibm.com/systems/power/hardware/blades/ibmi.html>

## How to get Redbooks

You can search for, view, or download Redbooks, Redpapers, Technotes, draft publications and Additional materials, as well as order hardcopy Redbooks, at this Web site:

[ibm.com/redbooks](http://ibm.com/redbooks)

## Help from IBM

IBM Support and downloads

[ibm.com/support](http://ibm.com/support)

IBM Global Services

[ibm.com/services](http://ibm.com/services)



# Index

## Symbols

/etc/dhcpd.conf 537  
/etc/init.d/dhcpd restart command 538  
/etc/init.d/dhcpd start command 538  
/etc/xinet.d/tftpd 540  
/sbin/service nfs reload command 544  
/sbin/service nfs start command 544

## Numerics

128-bit vector execution unit 46  
64-bit POWER6 processors 10

## A

active console 506  
Active Energy Manager (AEM) 10, 65  
Active Energy Manager (AEM) Console 408  
Active Memory Sharing (AMS) 177  
    configuration using IVM 183  
    DLPAR operations 208  
    planning considerations 179  
        consolidation factors 182  
        CPU resource sizing 181  
        memory weight 182  
        paging devices 181  
        prerequisites 179  
        workload 180  
    Shared Memory Pool. 180  
    summary 209  
add\_vlan\_ids 116  
Advanced Management Module 398  
Advanced Management Module (AMM) 251  
Advanced Management Web interface 64  
Advanced Peer-to-Peer Networking 312  
Advanced Program-to Program Communication (APPC) 312  
AIX 62  
AIX 6.1  
    Installation 211  
    LPAR Installation on VIOS 231  
AIX install kernel load 234  
Alert Standard Format (ASF) 16  
Alternate processor retry 43

AltiVec 46  
AltiVec (SIMD) 46  
AltiVec technology 46  
Anaconda 341  
Anchor Smartchip 50  
APPC 312  
APPN 312  
automated monitoring 66  
Automatic server restart (ASR) 18  
automatic service processor reset 18  
AutoYaST 521–522  
Autoyast 395  
AutoYaST main menu options 529  
    Hardware 530  
    Misc 532  
    Network Devices 530  
    Network Services 531  
    Security and Users 531  
    Software 529  
    System 530  
AutoYaST profile creation 522  
Available Virtual Disks 276

## B

Backup/Restore  
    IBM i 6.1 314  
Baseboard management controller (BMC) 16  
basic Kickstart configuration file 362  
Basic Zone Permission Table 255  
bit steering 23  
Blade Server features 16  
    Baseboard management controller (BMC) 16  
    Disk drive support 16  
    I/O expansion 17  
    IBM Director 17  
    Integrated network support 17  
    Large system memory capacity 17  
    Light path diagnostics 17  
    Power throttling 17  
    Reliability, Availability and Serviceability (RAS) 17  
BladeCenter  
    benefits 6

highlights 4  
BladeCenter H 26  
BladeCenter Hardware preparation 247  
BladeCenter JS23 18  
    Internal disk 21  
    Memory DIMMs 20  
    Memory features 20  
    Processor features 20  
BladeCenter JS43 21  
    Internal disk 23  
    Memory DIMMs 23  
    Memory features 23  
    Processor features 22  
BladeCenter Open Fabric Manager 7  
BladeCenter Power Domain 401  
boot device 83  
boot net command 396  
Bootstrap Protocol (BOOTP) 536  
bridge 224  
BRMS (Backup Recovery & Media Services) 314, 316

## C

cache availability 43  
Call Home contact message 483  
CDROM  
    virtual 228  
centralized BladeCenter environment 6  
CFFh 77  
CFFv 76  
chassis 24  
checkstop analysis 18  
chhwres 115  
Chipkill 23  
Chipkill technology 10  
chsyscfg 76  
chsysstate 170, 174  
client partition  
    create using IVM 218  
Cluster 1350™ portfolio 10  
Cluster Systems Management (CSM) 64, 66  
    Systems management  
        Cluster Systems Management 66  
CMM loan policy 181  
Collaborative Memory Manager (CMM) 181  
    policy 181  
Combined Form Factor horizontal 77  
Combined Form Factor vertical 76

commands  
AMM  
    env 405  
    fuelg 404  
chsyscfg 76  
IBM i  
    PWRDWNSYS 333  
    SHTDWNSYS 172  
install\_assist 239  
IVM  
    chsyscfg 99  
    chsysstate 170  
    LPAR shutdown 174  
lpcfgop 87  
lshwres 99  
lsparmigr 457  
lsrefcode 171  
lssyscfg, determine RMC status 442  
lssyscfg, determine system name 457  
lssyscfg, LPAR state 171  
lssyscfg, migrate status 457  
lsvet 436  
migrlpar 456  
license 87  
lsware 469  
lsmcode 469  
lssw 331  
migrlpar 457  
power 504  
printenv 511  
    printenv boot-device 511  
show-devs 511  
suma 63  
telnetcfg 504  
update\_flash 460  
validate 456  
VIOS  
    cfgassist 87–88  
    chdev, change hdisk reserve policy 439  
    entstat 76  
    help 85  
    ioslevel 435  
    license 87  
    lsdev 123  
    lsdev, check hdisk reserve policy 439  
    lspv 123  
    lstcpip 87, 90  
    mktcpip 87–88  
    mkvt 168

rmtcpip 93  
commit new firmware 460  
configuration file  
    Kickstart 362  
Connect the System i LAN console 280  
    Configuration wizard 282  
    dedicated service tools (DST) 289  
hosts file 285  
Service tool device ID 289  
System serial number 287  
Target partition 288  
Controlled end delay time 333  
Create Partition wizard 195  
Create Storage Pools 212  
Create virtual media library for backup  
    IBM i V6R1 Backup/Restore  
        Create virtual media library for backup 316  
creating an IBM i 6.1 partition 271  
Creating IBM i partition 272  
    Load Source and Console Identification 277  
current firmware level 469

**D**

Decimal floating point 44  
default ID 86  
Degraded boot support 18  
DHCP  
    AutoYaST configuration 531  
    configuration 537  
    restarting 538  
diagela 478  
diagnostics 66  
digital KVM 496  
Disk configuration in BladeCenter S 254  
    Basic Zone Permission Table 255  
    Individual BladeCenter S disk configuration 259  
    Predefined Config 258  
    Storage Configuration Manager (SCM) 566  
disk configuration in BladeCenter S  
    Zoning 255  
Disk consideration in BladeCenter S 252  
Disk Storage Modules (DSM) 253  
disk storage types 80  
Dynamic Host Configuration Protocol (DHCP) 536  
Dynamic logical partition (DLPAR) 208  
Dynamic Logical Partitioning (DLPAR) 365, 396, 546  
Dynamic LPAR (DLPAR) 156

capabilities, retrieve 157  
ethernet 161  
memory 158  
optical devices 164  
physical adapters 165  
processing 159  
storage 163  
Dynamic Power Saver Mode 47  
Dynamic Reconfiguration Tools package 443

## E

eConfigurator 10  
EnergyScale 49, 397  
Energyscale technology 47  
Enhanced SMT features 46  
entstat 76, 121  
error correction 23  
Ethernet bridge 224  
Ethernet I/O module 251  
Expansion cards 76  
    CFH  
        Qlogic combo card 77  
    CFFv  
        QLogic 4 Gbs fibre channel HBA 77  
    HSFF  
        Cisco 4x Infiniband 79  
Extended Error Handling (EEH) 18  
external storage 80

## F

fibre channel storage 80  
File Transfer Protocol (FTP)  
    installation server 536  
Firmware 493  
firmware file name 470  
firmware prompt  
    accessing 510  
firmware updates download for JS12/22 434  
First Failure Data Capture (FFDC) 6, 18  
Flexible Support Processor (FSP) 49  
from 462  
Front Panel (Control Panel) 473  
fuelg 404, 406

## G

get firmware level 469  
get the latest firmware image 471

Gigabit Ethernet 25

**H**

hdisk reserve policy 439  
help 85  
high performance computing (HPC) applications 5  
High Speed Form Factor 76  
HMC 271  
Host bus adapters (HBA) 83  
Host Ethernet Adapter (HEA) 56, 76, 100, 442  
    configuring as a SEA 103  
hotplug 546  
HSFF 76  
Hypertext Markup Language (HTML) 84  
Hypertext Transfer Protocol (HTTP) 536  
hypervisor 444

**I**

I/O hotplug 365, 396  
IBM BladeCenter chassis 24  
    BladeCenter H 25  
    BladeCenter HT 33  
    BladeCenter S 29, 80  
IBM Cluster Systems Management (CSM) 66  
IBM Director 64–65  
IBM Director benefits 39  
    Increased productivity of IT personnel and users 39  
    Reduced downtime 39  
    Reduced power consumption 39  
    Reduced service and support costs 39  
IBM i 64  
IBM i 6.1 64  
IBM i V6R1 Backup/Restore  
    IBM SAS LTO 314  
IBM i V6R1 install media preparation 280  
IBM i V6R1 installation 241  
    BladeCenter Hardware preparation 247  
    Disk configuration in BladeCenter S 254  
    Disk consideration in BladeCenter S 252  
    Hardware environments 244  
    installing PTFs 311  
    Preparation 242  
    Software installation process 242  
    Storage consideration BladeCenter H 252  
IBM i V6R1 installing PTFs  
    IBM ID 312  
IBM i V6R1 IPL types 292

IPL type A 292  
IPL type B 292  
IPL type C 292  
IPL type D 292  
IBM i V6R1 Navigator for i 312  
IBM i V6R1 Shutdown and restart 332  
    Shutdown IBM i V6R1 partition 332  
    Start an IBM i V6R1 partition 333  
IBM i V6R1 TCP/IP setup 312  
IBM i, storage considerations 82  
IBM ID 312  
IBM POWER6 processor 47  
IBM ServerProven 81, 565  
IBM service and productivity tools 396  
IBM service and productivity tools are packages 365  
IBM System Access for Windows V6R1 260  
    Installation System i Access for Windows 261  
    Preparation 260  
IBM Tivoli Storage Manager (TSM) 331  
    client check 331  
IBM Total Storage 80  
IBM Total Storage N Series 81  
IBM Workload Estimator (IBM WLE) 182  
iee\_virtual\_eth 116  
IEEE 1275 standard 509  
InfiniBand 25  
Install IBM i V6R1 operating system 303  
install\_assist 239  
installation  
    IBM AIX 6.1 212  
    IBM i 6.1 operating system 303  
    Red Hat Enterprise Linux 5.3 336  
    SUSE Linux Enterprise Server 11 368  
installation assistant 239  
installation of Red Hat Enterprise Linux 5.3 over the network 341  
installation of Red Hat Enterprise Linux 5.3 using a DVD 337  
installation of SLES 11 over the network 374  
installation of SLES 11 using a DVD 369  
installation summary  
    AIX 236  
Installation System i Access for Windows 261  
installing  
    Linux  
        configuring a BOOTP or DHCP service 537  
        network installation 536  
        virtual optical device setup and installation

544

Red Hat Enterprise Linux 5.3  
    automated installation 354  
System i Access for Window 261  
Installing License internal code (LIC) 292  
integrated Ethernet Switch module 251  
Integrated network support 17  
Integrated Virtual Ethernet (IVE) 56  
Integrated Virtualization Manager (IVM) 61, 374  
    command line interface 85  
    Create Partition button 145  
    Guided Setup view 94  
    Partition Properties 96  
        General tab 96  
        Memory tab 96  
        Processing tab 96  
    user interface 84, 194  
        navigation area 84  
        work area 84  
View/Modify Host Ethernet Adapters 101  
    Properties 101  
        Connected Partitions tab 102  
        General tab 101  
View/Modify Partitions link 95  
View/Modify Physical Adapters link 109  
View/Modify Virtual Ethernet, SEA configuration 105  
View/Modify Virtual Storage 122  
    Physical Volumes tab 122  
    Storage Pools tab 125  
    VIOS management 83  
Integrated Virtualization Manager V1.4 331  
Intelligent Copper Pass-Thru Module 76  
Intelligent Platform Management Interface (IPMI) 16  
Internal SAS 80  
interoperability matrix 82  
interposer card 79  
is\_trunk 116

## K

Kickstart configuration file 362  
Kickstart Configurator 356  
Kickstart installation 364  
Kickstart profile creation 354  
KVM 84, 496

## L

last known good firmware 460  
Licensed Internal Code (LIC) 292  
Light emitting diodes (LEDs) 18  
Light Path 475  
Light path diagnostics 17–18  
Linux 63  
Linux network installation 536  
live partition migration 365, 396, 546  
Live Partition Mobility 62, 81, 433  
    LPAR OS versions 437  
    Migrate option 453  
    migrate using CLI 457  
    migrate using IVM 453  
    migrating the LPAR 448  
        from the command line 456  
        using the IVM UI 448  
PowerVM Enterprise 436  
preparation 438  
preperation  
    Memory region size 438  
    networking 442  
    Storage and hdisk reserve policy 439  
requirements 434  
    firmware 434  
    hardware 434  
    VIOS version 434  
status using CLI 457  
status using IVM 454  
validate using CLI 456  
validate using IVM 448  
Logical over-commit 180  
logical partition (LPAR) 546  
    non-dynamic operations 84  
logical partitioning 61  
logical volume base 213  
LPAR activation  
    from the CLI 170  
    from the UI 168  
LPAR creation  
    Optical Devices 152  
    Partition Ethernet 149  
    Partition Memory 147  
    Partition Name 146  
    Partition Processors 148  
    Partition Storage Types and Assignments 150  
    Partition Summary 154  
    Physical Adapters 153  
LPAR shutdown 172

- from the CLI 174
  - from the UI 172
  - lpcfgop 87
  - lpcfgop command 87
  - lsparmigr 457
  - lsrefcode 171
  - lssyscfg 171, 442, 457
- M**
- machine checks 472
  - man dhcpcd command 537
  - media 84
    - Media Access Control (MAC) 537
    - media library 324
      - adding new media 137
      - create 135
      - create blank media 139
      - delete 137
      - extend 137
      - modify media assignment 141
    - media library size 324
    - media tray 84
  - memory
    - placement rules 51
    - memory controller 20, 23
    - memory region size 438
    - memory scrubbing 23
    - Memory subsystem 50
  - Micro-Partitioning 61
  - migrpar 456
  - mobile partition 455
  - monitoring
    - automated 66
    - resources 66
  - MSIM 79
    - Multi Switch Interconnect Module (MSIM) 79
  - Multiple Expansion Unit (MPE) 22
  - multi-processor expansion unit (MPE) 5
  - Multi-Switch Interconnect Module 79
- N**
- native Red Hat Enterprise Linux 5.2 installation 353
  - native SLES10 SP2 installation 394
  - Navigator for i 312
  - Network considerations
    - AMM 251
    - Ethernet I/O module 251
    - IBM i V6R1 installation 251
- PC for LAN Console 252
  - V6R1 LAN console 252
  - V6R1 production interface 252
  - VIOS/IVM 251
- Network File System (NFS) 341, 374, 536, 541
  - network installation 66
    - configuration file 537
    - installation source 536
    - Linux 536
    - Trivial File Transfer Protocol (TFTP) 539
  - network interface backup (NIB) 76
  - network interface controller (NIC) 500
  - NIM server 84
  - Non over-commit 180
  - Normal Mode Boot 343
  - Novell® SuSE Linux Enterprise Server 62
- O**
- Online Update and Package Selection 529
  - Open Firmware 493
    - boot parameters 364, 395
    - interface 509
  - Open Firmware prompt 509
  - Operating system support 62
    - IBM i 64
    - Linux 63
  - Optical devices 132, 213
    - physical 133
  - Optical Pass Thru Module 78
- P**
- paging devices 181
  - paging storage pool 184
  - partition
    - creating a partition
      - IBM i 6.1 271
    - creating partition
      - AIX 221
    - preparation 218
    - preparing PowerVM client partition 218
    - View/Modify Partitions page 219
  - Partition Management 221
  - Partition Mobility 76
  - PCI express 53
  - PCI-e 53
  - PCI-x 53
  - physical optical
    - remote media 133

physical optical device 215  
Physical over-commit 180  
physical volumes 122  
    assigning 226  
Point-to-point protocol (PPP) 16  
port\_vlan\_id 116  
POST 507  
Power cap 49  
Power Capping 48  
Power On Self Test 507  
Power Saver Mode 47  
Power Trending 47  
POWER® Hypervisor (PHYP) 472  
POWER6 444  
POWER6 Enhanced 444  
POWER6 Hypervisor (PHYP) 18  
POWER6 processor 43–44, 460  
    AltiVec (SIMD) 46  
    Decimal floating point 44  
    Energyscale technology 47  
    Simultaneous Multi Threading 45  
POWER6 processor based blade family 10  
POWER6+ 444  
POWER6+ Enhanced 444  
Power-on self-test (POST) 18  
PowerVM 61, 436  
    PowerVM editions for BladeCenter JS23 and JS43 62  
PowerVM editions for BladeCenter JS23 and JS43 62  
PowerVM key 436  
printenv 511  
    printenv boot-device 511  
Processor Core Nap 48  
Processor instruction retry 43  
PWRDWNSYS 333

## R

RAID SAS Switch Modules (RSSM) 59  
Red Hat Enterprise Linux 62  
    configuration tool 356  
    Kickstart 364  
Red Hat Enterprise Linux 5.2 automated installation 354  
Red Hat Enterprise Linux for POWER Version 4.6 336  
Red Hat Enterprise Linux for POWER Version 5.1 336

Redbooks Web site 568  
    Contact us xxxi  
reliability, availability and serviceability (RAS) 17  
reliability, availability, and serviceability (RAS) 365, 396, 546  
reliable double data rate 2 (DDR2) 5  
remote command execution 66  
remote control 496  
remote control Java applet 499  
Renesas Technology 15  
Reserved Firmware Memory 183  
Reserved Firmware Memory value 183  
reset-all 515  
resource balancing 46  
resource monitoring 66  
Resource Monitoring and Control (RMC) 442  
Restart Blade option 503  
rpm -Uvh command 550, 554

## S

SAS controller 58  
SAS Disk Array Manager for AIX 58  
SAS Modules 251  
SAS/SATA 80  
SCM 259  
secure shell (SSH) connection 500  
Serial Attached SCSI (SAS) 5, 10, 57  
Serial over LAN (SoL) 16, 84–85, 500  
ServerProven 81  
Service and Productivity Tools for Linux 545  
Service Processor 49  
Service Request Numbers (SRN) 479  
Service tool device ID 290  
service xinetd restart command 540  
set-connection-mode 515  
Shared Dedicated Capacity 62  
Shared Ethernet Adapter (SEA) 103, 442  
shared memory pool 208  
show-devs 511  
Shutdown IBM i V6R1 partition 332  
Simultaneous Multi Threading 45  
single processor checkstopping 43  
single rack-mounted chassis 4  
Single-threaded (ST) operation 46  
smit sasdamp 58  
SMS 493  
SMT 45  
SoL 84, 500

Solid State Disk Technology 5  
Solid state drive (SSD) 17  
specialized hardware detection circuits 472  
SSH 495  
SSIC 83  
Standard Form Factor 76  
Start an IBM i V6R1 partition 333  
Start Remote Control 498  
starting the firmware image  
    from the TEMP side 462  
Static Power Saver Mode 47  
storage area network (SAN) 341, 374  
Storage Configuration Manager (SCM) 259, 566  
Storage consideration BladeCenter H 252  
storage considerations 81  
Storage Management 212  
storage pool 184  
Storage Pools 124, 212  
    delete 126  
    reduce 126  
suma 63  
superuser rights 536  
support matrixes 81  
SUSE Linux Enterprise Server 62  
SUSE Linux Enterprise Server (SLES) 367  
SynapSense Wireless Sensor Node 66  
system access for Windows 260  
system diagnostics 472  
    diagnostic tools 472  
        checkpoints and error codes 472  
        for the AIX operating system 478  
        for the IBM i operating system 478  
        for the Linux operating system 478  
    Light Path and Front Panel diagnostics 473  
    stand-alone diagnostics 478  
IBM i partition diagnostics and errors 483  
reference codes 479  
    error codes 479  
    progress codes 479  
Service Advisor 483  
    using the Advanced Management Module 480  
System i Access for Windows 260  
System Management Service (SMS)  
    options 501  
System Reference Codes (SRCs) 479  
System specification  
    BladeCenter JS23 and JS43 chassis support  
    24  
    BladeCenter JS23 disk support 21  
BladeCenter JS23 memory support 20  
BladeCenter JS23 processor support 20  
BladeCenter JS43 disk support 23  
BladeCenter JS43 memory support 23  
BladeCenter JS43 processor support 23  
BladeCenter support 25  
JS23 standard configuration 19  
JS43 configuration options 21  
System Storage Interoperation Center 83  
system-config-kickstart command 356  
Systems management 64  
    Advanced Management Web interface 64  
    IBM Director 65

**T**

telnet 495  
telnetcfg 504  
tftpboot 539  
Thermal Power Management Device (TPMD) 49  
    POWER6 processor 49  
thread priority 46  
Trivial File Transfer Protocol (TFTP) 536  
    configuration 539  
    network installation 536  
trunk\_priority 116  
TSM client check 331

**U**

Universal Serial Bus (USB) subsystem 55  
Unmount All option 330  
USB 55

**V**

Very Low Profile (VLP) RDIMM packages 20  
VIOS 177  
    planning 179  
VIOS/IVM 251  
Virtual CDROM 228  
virtual disks 127  
    assigning 226  
    create 128  
    delete 131  
    extend 129  
Virtual Ethernet  
    Integrated Virtualization Manager (IVM) 84  
Virtual Ethernet Adapters 103  
Virtual I/O Server (VIOS) 61, 71

data sheet 81, 565  
default user 86, 94  
fixpack download location 434  
planning 73  
Virtual Input/Output Server 177  
virtual media library 212  
virtual optical device 152, 544  
virtual optical devices 447  
Virtual Storage Management 121  
virtual tape 152  
Virtual Terminal 167  
Vital product data (VPD) 18  
VMLibrary 216

## **W**

Windows 260  
WRKOPTVOL 321

## **Y**

yast command 539  
YaST graphical interface 523  
yast2 command 523

## **Z**

Zone Group 255  
Zoning 255





**Redbooks**

# IBM BladeCenter JS23 and JS43 Implementation Guide

(1.0" spine)  
0.875" <-> 1.498"  
460 <-> 788 pages





# IBM BladeCenter JS23 and JS43 Implementation Guide



**Featuring installation  
techniques for the  
IBM AIX, IBM i, and  
Linux**

**Showing Live  
Partition Mobility  
scenarios**

**Detailed coverage of  
AMS, IVM and power  
management**

This IBM Redbooks publication provides a detailed technical guide for configuring and using the IBM BladeCenter JS23 and IBM BladeCenter JS43 servers. These IBM Power Blade servers feature the latest IBM POWER6 processor technology. This book teaches you how to set up the latest Power Blade servers to run AIX, i, and Linux operating systems in the IBM BladeCenter architecture.

This book will be useful for data center system architects; network, storage and facilities engineers; system administrators; and application architects that need a detailed system bring-up guide and integration guide for IBM Power Blade servers hosting virtualized operating system partitions. We include procedures for installation and configuration of Virtual Input/Output Server (VIOS), Integrated Virtualization Manager (IVM), IBM AIX Version 6.1, IBM i 6.1 (i5/OS V6R1), Red Hat Enterprise Linux, and SUSE Enterprise Linux.

**INTERNATIONAL  
TECHNICAL  
SUPPORT  
ORGANIZATION**

**BUILDING TECHNICAL  
INFORMATION BASED ON  
PRACTICAL EXPERIENCE**

IBM Redbooks are developed by the IBM International Technical Support Organization. Experts from IBM, Customers and Partners from around the world create timely technical information based on realistic scenarios. Specific recommendations are provided to help you implement IT solutions more effectively in your environment.

**For more information:  
[ibm.com/redbooks](http://ibm.com/redbooks)**